EXPERT INSIGHT

AI Product Manager's Handbook

The ultimate playbook to unlock AI product success with real-world insights and strategies

Second Edition

Irene Bratsis

packt

AI Product Manager's Handbook

Second Edition

The ultimate playbook to unlock AI product success with real-world insights and strategies

Irene Bratsis

<packt>

AI Product Manager's Handbook

Second Edition

Senior Publishing Product Manager: Tushar Gupta
Acquisition Editor – Peer Reviews: Swaroop Singh
Project Editor: Namrata Katare
Content Development Editor: Shruti Menon
Copy Editor: Safis Editing
Technical Editor: Kushal Sharma
Proofreader: Safis Editing
Indexer: Hemangini Bari
Presentation Designer: Ajay Patule
Developer Relations Marketing Executive: Vignesh Raju

First published: February 2023
Second edition: November 2024

Production reference: 5070725

Published by Packt Publishing Ltd.
Grosvenor House
11 St Paul's Square
Birmingham
B3 1RB, UK.

ISBN 978-1-83588-284-9

`www.packt.com`

For anyone who's ever been lost in the woods. For those who find me over and over again.

– Irene Bratsis

Contributors

About the author

Irene Bratsis currently oversees product at Presage, a production company and magazine based in NYC, works independently as a consultant, and most recently served as Director of Digital Product and Data at the International WELL Building Institute (IWBI). She has a bachelor's degree in economics and international relations from Simmons University. Irene has completed a big data and analytics certificate from MIT xPRO and an intensive data science program with Thinkful.

Before IWBI, she worked as an operations analyst at Tesla, a data scientist at Gesture, a data product manager at Beekin, and led product at Tenacity. Irene has served as the co-lead of the NYC chapter of Women in Data and currently organizes and moderates WaiTalks, an online series of talks organized by Women in AI. A founding member of Trusted AI, Irene has written articles for well-known outlets such as Data Driven Investor, Analytics Vidhya, and Towards Data Science, as well as her own Medium and LinkedIn profiles. Irene is currently based in NYC.

About the reviewer

Josh Atlas is an accomplished product management leader with over 15 years of experience, spanning fintech, e-commerce, AI/ML, and digital marketing. He has a proven track record of driving product innovation and operational efficiencies, delivering impactful solutions for global organizations like Meta, Google Nest, and Walmart, as well as startups such as Vivid Labs. His expertise includes leading cross-functional teams, navigating complex challenges, and building customer-centric products, including AI-driven tools and next-generation NFT applications.

In addition to his work in traditional product management, Josh has designed and built advanced AI agents, including Josh-GPT, a digital twin that serves as an interactive extension of his professional knowledge and experience and reflects his deep commitment to innovation in generative AI and conversational technologies. His contributions extend to transformative projects, such as reducing loan manufacturing time at GoodLeap, launching a one-click badge creation system at Meta, and designing a customer-focused AI chatbot for solar energy adoption.

Josh is passionate about leveraging data, collaboration, and design thinking to create products that solve real-world problems while fostering strategic alignment across teams and stakeholders.

I would like to express my heartfelt gratitude to my wife for her patience and unwavering support, which have been the foundation of my success. Your encouragement and belief in me have carried me through every challenge and achievement. To my family and friends, thank you for your constant inspiration and love – I could not have done this without you.

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

Table of Contents

Preface

It's hard to find anyone these days who doesn't have strong reactions to AI. I've watched my own feelings evolve with its rise, ebbing and flowing over the years. As a student, I felt an overwhelming excitement and optimism about where AI – and the fourth industrial revolution it accompanies – might lead us. That initial thrill was tempered as I began organizing AI events with virtual speakers and managing a data and AI book club. I adopted a monthly practice of learning about how bias and dependence on AI compromise our lives in both visible and unseen ways. AI is a double-edged sword – capable of driving immense progress but fraught with ethical dilemmas, privacy risks, and the perpetuation of biases we're still struggling to confront in the real world today.

And so, we arrive at one of the greatest debates that resurfaces with every technological leap: do we dare embrace powerful technology even when we're aware of the risks? As far as I see it, we don't really have a choice – the debate itself is an illusion we indulge in. With the rise of accessible, generative AI tools available today, it's clear that it's here to stay. Nihilistic fears about it won't protect us from harm. Pandora's box is open, and as we peer into what remains inside, we find that hope springs eternal. AI is shaping our future, whether we're ready or not. It has the potential to enhance human creativity and address pressing global challenges. Yet, the more we integrate this technology, the more we must ensure that AI serves humanity, and not just the interests of a few. Philosophically, the questions AI raises about intelligence and consciousness are essential to redefining what it means to be human in an age where machines can think, adapt, and even create.

I wanted to write a book about AI product management because it's the makers of products who transform possibilities into realities. Understanding the intricacies of how to ideate, build, manage, and sustain AI products with integrity, to the best of my ability, feels like the greatest contribution I can offer to this field at this moment in time. I'm encouraged by the collective bargaining power of individuals demanding that companies adopt AI ethically and responsibly. I'm relieved that so many AI product teams today prioritize human-centered design and are committed to building products they can proudly bring to market. This shift holds a mirror to our biases and prejudices, prompting us to look deeply into the reflection we see – asking whether we truly like what we've created. It places the human experience of AI front and center, encouraging us to build expressions of AI that reflect our highest aspirations rather than our deepest flaws.

It's been an honor to deliver this second edition.

Who this book is for

This book is for people that aspire to be AI product managers, AI technologists, and entrepreneurs, or for people that are casually interested in the considerations of bringing AI products to life. It should serve you if you're already working in product management and you have a curiosity about building AI products. It should also serve you if you already work in AI development in some capacity and you're looking to bring those concepts into the discipline of product management and adopt a more business-oriented role. While some chapters in the book are more technically focused, all of the technical content in the book can be considered beginner level and accessible to all.

Part 1 of this book is meant to serve as an overview of topics spanning the AI landscape overall, types of product that can exist in the space and a glance at the industry as a whole. *Part 2* will have more practical, applied content regarding the product management of AI native tools. *Part 3* will keep this format, but will focus on transitioning a traditional software product into an AI product. In these two parts, you'll see more diagrams, flow charts, checklists, and visual aids suitable for a handbook. Finally, *Part 4*, the newest part of the book, will focus on the management of an AI career itself, serving as a handbook for maturing in the PM role and the pathways you can take with it.

What this book covers

Chapter 1, *Understanding the Infrastructure and Tools for Building AI Products*, offers an overview of the main concepts and areas of infrastructure for managing AI products.

Chapter 2, *Model Development and Maintenance for AI Products*, delves into the nuances of model development and maintenance.

Chapter 3, *Machine Learning and Deep Learning Deep Dive*, is a broader discussion of the difference between traditional deep learning and deep learning algorithms and their use cases.

Chapter 4, *Commercializing AI Products*, discusses the major areas of AI products we see in the market, as well as examples of the ethics and success factors that contribute to commercialization.

Chapter 5, *AI Transformation and Its Impact on Product Management*, explores the ways AI can be incorporated into the major market sectors in the future.

Chapter 6, *Understanding the AI-Native Product*, provides an overview of the strategies, processes, and team building needed to empower the success of an AI-native product.

Chapter 7, *Productizing the ML Service*, is an exploration of the trials and tribulations that may come up when building an AI product from scratch.

Chapter 8, *Customization for Verticals, Customers, and Peer Groups*, is a discussion on how AI products change and evolve over various types of verticals, customer types, and peer groups.

Chapter 9, *Product Design for the AI-Native Product*, is an overview of product design principles and concepts that are customized for products built natively with AI/ML components.

Chapter 10, *Benchmarking Performance, Growth Hacking, and Cost*, explains the benchmarking needed to gauge product success at the product level rather than the model performance level.

Chapter 11, *Managing the AI-Native Product*, reviews ongoing AI PM considerations that relate to leadership and visionary, stakeholder and operational alignment of products built natively with AI.

Chapter 12, *The Rising Tide of AI*, revisits the concept of the Fourth Industrial Revolution and a blueprint for products that don't currently leverage AI.

Chapter 13, *Trends and Insights across Industry*, dives into the various ways we're seeing AI trending across industries, as well as accessible routes product teams can take when enabling AI

Chapter 14, *Evolving Products into AI Products*, is a practical guide on how to deliver AI features and upgrade the existing logic of products to successfully update products for AI commercial success.

Chapter 15, *The Role of AI Product Design*, refocuses AI design and communication foundations applied to product teams that are looking to evolve traditional software products with AI/ML capabilities.

Chapter 16, *Managing the Evolving AI Product*, reviews ongoing AI PM considerations that relate to leadership and visionary, stakeholder and operational alignment of traditional software products adopting AI features and capabilities.

Chapter 17, *Starting a Career as an AI PM*, brings readers striving for AI PM careers on a journey through the theoretical and applied foundations to set up their budding careers up for success.

Chapter 18, *What Does It Mean to Be a Good AI PM?*, breaks down the various facets of an AI PM and the technical, business, communication, leadership and problem solving considerations for those looking to excel in the role.

Chapter 19, *Maturing and Growing as an AI PM*, explores the various ways AI PMs can mature in their careers through projecting their ideal AI PM roadmap, staying informed with learning paths, networking to deepen connections and sharing their experiences and wisdom with others.

Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here: `https://packt.link/gbp/9781835882849`.

Conventions used

There are a number of text conventions used throughout this book.

`CodeInText`: Indicates code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. For example: "One of our hyperparameters for this random forest example was setting our cross-validation to `10`.

Bold: Indicates a new term, an important word, or words that you see on the screen. For instance, words in menus or dialog boxes appear in the text like this. For example: " This phenomenon is called **overfitting** and it's a big topic of conversation in data science and ML circles."

Warnings or important notes appear like this.

Tips and tricks appear like this.

Get in touch

Feedback from our readers is always welcome.

General feedback: Email `feedback@packtpub.com` and mention the book's title in the subject of your message. If you have questions about any aspect of this book, please email us at `questions@packtpub.com`.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you reported this to us. Please visit `http://www.packtpub.com/submit-errata`, click **Submit Errata**, and fill in the form.

Piracy: If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at `copyright@packtpub.com` with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit `http://authors.packtpub.com`.

Leave a Review!

Thank you for purchasing this book from Packt Publishing—we hope you enjoy it! Your feedback is invaluable and helps us improve and grow. Once you've completed reading it, please take a moment to leave an Amazon review; it will only take a minute, but it makes a big difference for readers like you.

`https://packt.link/r/1835882854`

Part 1

Lay of the Land – Terms, Infrastructure, Types of AI, and Products Done Well

An AI product manager needs to have a comprehensive understanding of AI, along with all the varied components that lead to its success, if they're going to be successful in commercializing their products. This first part consists of five chapters that will cover what the term AI encompasses and how to support infrastructure to make it successful within your organization. It will also cover how to support your AI program from a maintenance perspective, navigate the vast areas of **machine learning** (**ML**) and **deep learning** (**DL**), choose the best path for your product, and understand current and future developments in AI products.

By the end of this part, you will understand AI terms and components, what AI implementation means from an investment perspective, how to maintain AI products sustainably, and how to choose between the types of AI that would best fit your product and market. You will also learn about the success factors for ideating and building a **minimum viable product** (**MVP**) and how to make a product that truly serves its market.

This part comprises the following chapters:

- *Chapter 1*, *Understanding the Infrastructure and Tools for Building AI Products*
- *Chapter 2*, *Model Development and Maintenance for AI Products*
- *Chapter 3*, *Deep Learning Deep Dive*
- *Chapter 4*, *Commercializing AI Products*
- *Chapter 5, AI Transformation and Its Impact on Product Management*

1

Understanding the Infrastructure and Tools for Building AI Products

The frontier of **artificial intelligence** (**AI**) products seems a lot like our universe: ever-expanding. That rate of expansion is increasing with every passing year as we go deeper into a new way to conceptualize the products, organizations, and industries we're all a part of. Laying a solid foundation is an essential part of understanding this transformation, which is our goal with this book. Since virtually every aspect of our lives is expected to be impacted in some way by AI, we hope you will come out of this experience more confident about what AI adoption will look like for the products you support or hope to build someday.

Part 1 of this book will serve as an overview of the lay of the land. We will cover terms, infrastructure, types of AI algorithms, and products done well, and by the end of this part, you will understand the various considerations when attempting to build an AI strategy, whether you're looking to create a native-AI product or add AI features to an existing product. We will be covering theoretical concepts in *Part 1* and will be using *Parts 2* and *3* for more practical content, where we will see these theoretical concepts applied more specifically. In *Part 4*, we will discuss how you can build and grow your AI PM career.

Managing AI products is a highly iterative process, and the work of a product manager (PM) is to help your organization discover what the best combination of infrastructure, training, and deployment workflow is to maximize success in your target market. The performance and success of AI products lie in understanding the infrastructure needed for managing AI pipelines, the outputs of which will then be integrated into a product. In this chapter, we will cover everything from databases to deployment strategies to tools you can use to manage your AI projects, as well as how to gauge your product's efficacy.

This chapter will serve as a high-level overview of the subsequent chapters in *Part 1*, but it will first and foremost provide a definition of terms, which is often quite hard to come by in today's marketing-heavy AI competitive landscape. These days, it feels like every product is an AI product, and marketing departments are trigger-happy with sprinkling that term around, rendering it almost useless. A big part of why I wanted to write this book is because of my own challenges with navigating competitive landscapes in hopes of trying to understand various "AI" products.

I felt that if I had trouble understanding how products out there were built when using AI, given my familiarity with machine learning, what hope could there be for others? I suspect this won't be changing anytime soon, but the more fluency consumers and customers alike have with the capabilities and specifics of AI, **machine learning** (**ML**), **deep learning** (**DL**), and data science, the more we should see clarity about how products are built and optimized. Understanding the context of AI is important for anyone considering building or supporting an AI product.

In this chapter, we will cover the following topics:

- Definitions – what AI is and is not
- The old – exploring ML
- The new – exploring DL
- ML versus DL – understanding the difference
- Learning paradigms in ML
- The order – what is the optimal flow and where does every part of the process live?
- DB 101 – databases, warehouses, data lakes, and lakehouses
- Managing projects – IaaS
- Deployment strategies – what do we do with these outputs?
- Succeeding in AI – how well-managed AI companies do infrastructure right
- The promise of AI – where is AI taking us?

Getting the most out of this book – get to know your free benefits

Unlock exclusive **free** benefits that come with your purchase, thoughtfully crafted to supercharge your learning journey and help you learn without limits.

Here's a quick overview of what you get with this book:

Next-gen reader

Figure 1.1: Illustration of the next-gen Packt Reader's features

Our web-based reader, designed to help you learn effectively, comes with the following features:

Multi-device progress sync: Learn from any device with seamless progress sync.

Highlighting and notetaking: Turn your reading into lasting knowledge.

Bookmarking: Revisit your most important learnings anytime.

Dark mode: Focus with minimal eye strain by switching to dark or sepia mode.

Interactive AI assistant (beta)

Figure 1.2: Illustration of Packt's AI assistant

Our interactive AI assistant has been trained on the content of this book, so it can help you out if you encounter any issues. It comes with the following features:

✦ **Summarize it:** Summarize key sections or an entire chapter.

✦ **AI code explainers:** In the next-gen Packt Reader, click the **Explain** button above each code block for AI-powered code explanations.

Note: The AI assistant is part of next-gen Packt Reader and is still in beta.

DRM-free PDF or ePub version

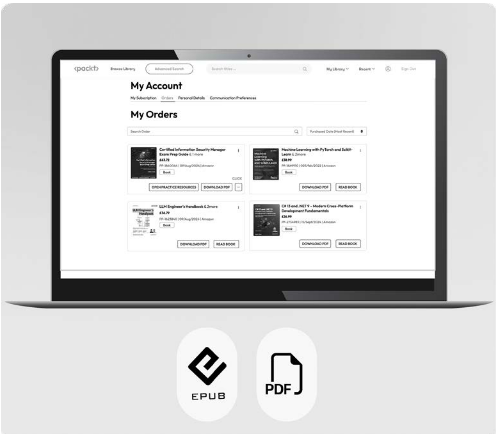

Figure 1.3: Free PDF and ePub

Learn without limits with the following perks included with your purchase:

Learn from anywhere with a DRM-free PDF copy of this book.

Use your favorite e-reader to learn using a DRM-free ePub version of this book.

Unlock this book's exclusive benefits now

Take a moment to get the most out of your purchase and enjoy the complete learning experience.

Note: Have your purchase invoice ready before you begin.

UNLOCK NOW

`https://www.packtpub.com/unlock/9781835882849`

Definitions – what AI is and is not

In 1950, a mathematician and World War II war hero named Alan Turing asked a simple question in his paper *Computing Machinery and Intelligence*, and that question was, *Can machines think?* Today, we're still grappling with that same question. Perhaps more so now that we have accessible, powerful **large language models** (**LLMs**) like ChatGPT and Claude to blur the lines. Depending on who you ask, AI can be many things. Many maps exist on the internet to define the broad categories of AI, from expert systems used in healthcare and finance to simpler forms of ML to more advanced models like neural networks. As we continue with this chapter, we will cover many of these facets of AI, particularly those that apply to products emerging in the market today. For the purposes of applied AI in products across industries, in this book, we will focus primarily on various applications of ML and DL models because these are often used in production anywhere AI is referenced in any marketing capacity. We will use AI/ML as a blanket term covering a span of ML applications, and we will cover the major areas most people would consider ML, such as DL, computer vision, natural language processing, and facial recognition. These are the methods of applied AI that most people will come across in the industry, and familiarity with these applications will serve any PM looking to break into AI. If anything, we'd like to help anyone who's looking to expand into the field from another product management background to choose which area of AI appeals to them most.

First, let's look at what is and what isn't ML. The best way for us to express it as simply as we can is: if a machine is learning from some past behavior and its success rate is improving as a result of this learning, it is ML! *Learning is the active element*. No models are perfect, but we do learn a lot from employing models. Most models will have some element of hyperparameter tuning, but all models will have parameters that are learned during the training process. The use of each model will yield certain results in performance, which data scientists and ML engineers will use to benchmark performance and improve upon it.

If there are fixed, hardcoded rules that don't change, it's not ML. AI is a field of computer science, and all programmers are effectively doing just that: giving computers a set of instructions to fire away on. If your current program doesn't learn from the past in any way and it simply executes on directives it was hardcoded with, we can't call this ML. Rather, you may have heard the terms *rules-based engine* or *expert system* thrown around in other programs. They are considered forms of AI but they're not ML because, although they are a form of AI, the rules are effectively replicating the work of a person, and the system itself is not learning or changing on its own.

We find ourselves in a tricky time in AI adoption where it can be very difficult to find information online about what makes a product AI. Marketing is eager to add the AI label to their products but there still isn't a baseline that explains what that means out in the market. This further confuses the term AI for consumers and technologists alike. If you're confused by the terms, particularly when they're applied to products you see promoted online, you're very much not alone.

Another area of confusion is the general term that is *AI*. For most people, the concept of AI brings to mind the *Terminator* franchise from the 1980s, the 2013 film *Her* of an OS gone wrong, and other futurist depictions of inescapable technological destruction. While there certainly can be a lot of harm to come from the AI we use today, the depictions in the films above represent what's referred to as *strong AI* or **artificial general intelligence** (**AGI**). AGI is a theorized, potential future state of AI where machines can perform high-value economic tasks without any human oversight or intervention. We still have ways to go for something as advanced as AGI but we've got plenty of what's referred to as **artificial narrow intelligence** or **narrow AI** (**ANI**).

ANI is also commonly expressed as *weak AI* and is what's generally meant when you see the term "AI" plastered all over products you find online. ANI is exactly what it sounds like: a narrow application of AI. Maybe it's good at talking to you, at predicting some future value, or at organizing things; maybe it's an expert at that, but its expertise won't bleed into other areas. For instance, GPT-4, a model that powers one of the most widely recognized conversational AIs, ChatGPT, might have advanced comprehension and conversation skills but it can't perform spinal surgery. If it could, it would stop being ANI. These major areas of AI are referred to as "strong" and "weak" in comparison to human intelligence.

But even that may be a controversial opinion. The Microsoft Research team published a paper in April 2023 detailing their exploration into GPT-4 and its proximity to AGI, titled *Sparks of artificial general intelligence: Early experiments with GPT-4*. The following is an excerpt from their abstract:

"We demonstrate that, beyond its mastery of language, GPT-4 can solve novel and difficult tasks that span mathematics, coding, vision, medicine, law, psychology and more, without needing any special prompting. Moreover, in all of these tasks, GPT-4's performance is strikingly close to human-level performance, and often vastly surpasses prior models such as ChatGPT. Given the breadth and depth of GPT-4's capabilities, we believe that it could reasonably be viewed as an early (yet still incomplete) version of an artificial general intelligence (AGI) system."

For every person out there who's come across Reddit threads about AI being sentient or somehow having ill will toward us, we want to make the following statement very clear. A complete form of AGI does not yet exist, and even if it does, we have no reason to believe it to be a sentient AI. This does not mean AI doesn't actively and routinely cause humans harm, even in its current forms. The major caveat here is that unethical, haphazard applications of AI already actively cause us both minor inconveniences and major upsets. Books like Cathy O'Neil's *Weapons of Math Destruction* detail the risks that come from AI using a number of real-world examples. Building AI ethically and responsibly is still a work in progress, and it's one that anyone who's involved in strategy, operations, and leadership should take seriously. We all have a role to play in building equitable applications of AI today and in the future. While AI systems may not be sentiently plotting the downfall of humanity, they inadvertently do cause harm when they're left untested, improperly managed, and inadequately vetted for bias.

For now, can machines walk and talk like us? Increasingly so, every day. Do they think like us? LLMs can demonstrate step-by-step thought processes that mirror those of humans. But are they actually sentient? Sentience has to do with the ability to experience feelings, emotions, and consciousness. Algorithms and data are powerful, but not powerful enough to draw blood from a stone. AI doesn't experience the world like us. It doesn't experience the neurochemical reactions we do. It won't internalize the gripping shame of abandonment or lose someone it loves to the opioid crisis. It doesn't fall in love, experience depression, worry about keeping food on the table, or fear homelessness. It's my personal opinion that the insufferable aspects of the human condition end with us. But I do very much believe that some of our greatest struggles, as well as our wildest curiosities, will be impacted considerably by the benevolence of AI and ML.Getting the most out of this book – get to know your free benefits

Introducing ML and DL

We have discussed how we've grappled with the idea of using machines since the 1950s, but we want to expand on the history of ML and DL **artificial neural networks** (**ANNs**) to give you a sense of how long these models have been around in order to give greater context and demonstrate the evolution these technologies have experienced to date.

The old – exploring ML

ML models attempt to create some representation of reality in order to help us make some sort of data-driven decision. Essentially, we use mathematics to represent some phenomenon that's happening in the real world. ML essentially takes mathematics and statistics to predict or classify some future state. The paths diverge in one of two ways:

- The first group lies with the emergence of models that continue to progress through statistical models.
- The second group lies with the emergence of models that try to mimic our own natural neural intelligence.

Colloquially, these are referred to as traditional ML and DL models, respectively. In this section, we will take a look at the traditional statistical ML models in order to understand both the historical relevance and prevalence of ML models. Some of the most reliable and prevalent models used in ML have been around for ages. **Linear regression** models have been around since the late 1800s and were popularized through the work of Karl Pearson and Sir Francis Galton, two British mathematicians. Their contributions gave way to one of the most popular ML algorithms used today, although unfortunately, both were prominent eugenicists. Karl Pearson is also credited with inventing **principal component analysis** (**PCA**), a learning method that reduces dimensions in a dataset, in 1901.

> In the context of datasets, dimensions are categories or variables you would use to describe and analyze your data. Think of them as "labels" or "attributes" that give context to the elements you have within a dataset. For example, a customer database might include the age, gender, and location dimensions

A popular ML method, **naive Bayes classifiers**, came onto the scene in the 1960s but they're based on the work of an English statistician named Thomas Bayes and his theorem of conditional probabilities, which is from the 1700s. The logistic function was introduced by Belgian mathematician Pierre Francois Velhulst in the mid-1800s, and **logistic regression** models were popularized by a British statistician named David Cox in 1958.

One of the simplest ML models for classification and regression, the **KNN algorithm**, emerged from a technical analysis report that was done by statisticians Evelyn Fix and Joseph Lawson Hodges Jr. on behalf of the United States Armed Forces in collaboration with Berkeley University in 1951. K-means clustering, an ML clustering method, was first proposed by a mathematician at UCLA named James MacQueen in 1967. As you can see, many of the algorithms that are used most commonly in ML models today have their roots quite far back in our modern history. Their simplicity and elegance add to their relevance today.

A brief history of DL

In 1943, Warren S. McCulloch and Walter Pitts published a paper, *A logical calculus of the ideas immanent in nervous activity*, which made a link between mathematics and neurology by creating a computer model based on the neural networks inherent in our own brains based on a combination of algorithms to create a "threshold" to mimic how we pass information from our own biological network of neurons. Then, in 1958, Frank Rosenblatt published a paper that would be widely considered the ancestor of neural nets, called *The Perceptron: A perceiving and recognizing automaton*. This was, for all intents and purposes, the first, simplest, and oldest ANN.

In the 1960s, developments toward backpropagation, or the idea that a model learns from layers of past mistakes as it trains its way through a dataset, made significant strides toward what would eventually make up the neural network. The most significant part of the development that was happening at this time was coupling the idea of inspiring mathematical models with the way the brain works based on networks of neurons and backpropagation because this created the foundation of ANNs, which learned through past iterations.

It's important to note here that many ANNs work in a "**feedforward**" motion in that they go through the input, hidden layers, and output layers sequentially and in one direction only, from input to output. The idea of **backpropagation** essentially allows the ANNs to learn bi-directionally so that they're able to minimize the error in each node, resulting in better performance. The following diagram illustrates this:

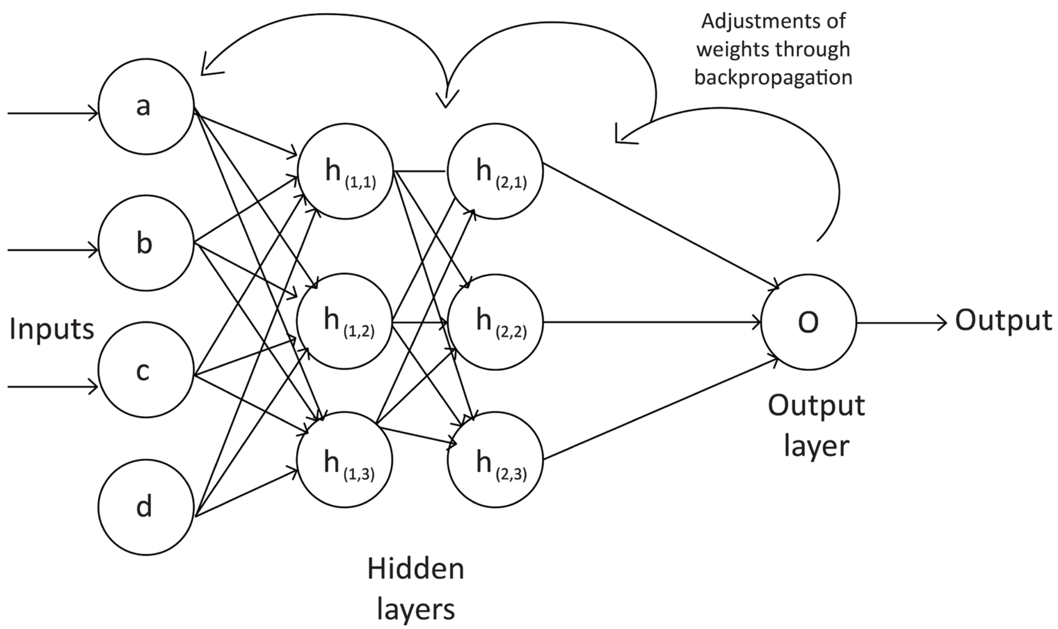

Figure 1.4: Backpropagation

It wasn't until 1986, when David Rumelhart, Geoffrey Hinton, and Ronald Williams published a famous paper, *Learning representations by back-propagating errors*, that people fully began to appreciate the role backpropagation plays in the success of DL. The idea that you could backpropagate through time, allowing neural networks to assign the appropriate weights as well as train a neural network with hidden layers, was revolutionary at the time.

After each development, there was much excitement for ML and the power of neural networks but between the mid-60s and the 80s, there were two significant issues: a lack of data as well as a lack of funding. If you've heard the term "AI winter," this is what it's referring to. Developments were made on the modeling side but we didn't have significant ways to apply the models that were being developed without the ability and willingness of research groups to get their hands on enough data to feed those models.

Then, in 1997, Sepp Hochreiter and Jürgen Schmidhuber published their groundbreaking work titled *Long Short-Term Memory*, which effectively allowed DL to "solve complex, artificial long-time lag tasks that had never been solved by previous recurrent network algorithms." The reason why this development was so important was it allowed the idea of sequences to remain relevant for DL problems. Because neural networks involve hidden layers, it's difficult for the notion of time to remain relevant, which makes a number of problems hard to solve. For instance, a traditional recurrent neural network might not be able to autocomplete a sentence in the way that a **long short-term memory (LSTM)** can because it doesn't understand the time sequence involved in the completion of a sentence.

Today, most DL models require a ton of data, meaning the neural networks that power DL need lots of examples to understand whether something is, for example, a dog or a horse. If you think about it a bit, though, this doesn't actually relate that closely to how our brains work. A small child that's just emerging and learning about the world might need to be reminded once or twice about the difference between a dog and a horse, but you likely aren't reminding them of that difference thousands or millions of times. In that sense, DL is evolving toward requiring fewer and fewer examples to learn. Sure, these days, we're able to gather massive amounts of data for DL models to learn from, but the models themselves are evolving to improve without needing much data toward the ultimate goal of DL models that can be trained with small amounts of data.

So far, we've covered some of the history and influences shaping the field of ML and DL more specifically. While we haven't gone into many of the technical concepts too heavily, this gives us a good foundation with which to understand how ML and DL have developed over time and why they've risen to prominence. In the following section, we will get more hands-on and get into understanding DL better.

The new – exploring DL

Part of our intention with separating ML and DL conceptually in this book is really to create associations in your mind regarding these concepts. For most technical folks in the field, there are specific models and algorithms that come up when you see ML or DL as a descriptor of a product. Quick reminder here that DL is a subset of ML. If you ever get confused by the two terms, just remember that DL is a form of ML that's grown and evolved to form its own ecosystem. Our aim is to demystify that ecosystem as much as possible so that you can confidently understand the dynamics at play with DL products as a PM.

The foundational idea of DL is centered around our own biological neural networks, and DL uses what's often the umbrella term of ANNs to solve complex problems. As we will see in the next section, much of the ecosystem that's been formed in DL has been inspired by our own brains, where the "original" neural networks are found. This inspiration comes not just from the function of the human brain, particularly the idea of learning through examples, but also from its structure.

Because this isn't an overly technical book meant for DL engineers, we will refrain from going into the terms and mathematics associated with DL. A basic understanding of an ANN would be helpful, however. As we go through this section, keep in mind that a neural network is composed of artificial neurons or nodes and that these nodes are stacked next to one another in layers. Typically, there are three types of layers:

- The input layer
- The hidden layer(s)
- The output layer

While we will go over the various types of ANNs, there are some basics to how these DL algorithms work. Think in terms of layers and nodes. Essentially, data is passed through each node of each layer and the basic idea is that there are weights and biases that are passed from each node and layer. The ANNs work through the data they're training on in order to best arrive at patterns that will help them solve the problem at hand. An ANN that has at least three layers (which means an input, an output, and a minimum of one hidden layer) is "deep" enough to be classed as a DL algorithm.

That settles the layers, but what about the nodes? One of the simplest models (which we will discuss in detail later in this chapter) is the linear regression model. You can think of each node as its own mini-linear regression model because this is the calculation that's happening within each node of an ANN. Beyond that, each node does more than just a linear regression calculation as it also applies a non-linear activation function to the result. This activation function then introduces non-linearity into the model, allowing ANNs to capture complex patterns and relationships in the data that linear regression cannot. Each node has its data, a weight for that data, and a bias or parameter that it's measuring against to arrive at an output. The summation of all these nodes making these calculations at scale gives you a sense of how an ANN works. If you can imagine a large scale of hundreds of layers, each with many nodes within each layer, you can start to imagine why it can be hard to understand why an ANN arrives at certain conclusions.

DL is often referred to as a black-box technology and this starts to get to the heart of why that is. Depending on our math skills, we humans can explain why a certain error rate or loss function is present in a simple linear regression model. We can conceptualize the ways a model, which is being fitted to a curve, can be wrong. We can also appreciate the challenge when presented with real-world data, which doesn't lay out a perfect curve, for a model. But if we increase that scale and try to conceptualize potentially billions of nodes, each representing a linear regression model, our brains will start to hurt.

Though DL is often discussed as a bleeding-edge technological advancement, as we saw in the prior section, this journey also started long ago.

Invisible influences

It's important to understand the underlying relationships that have influenced ML and DL as well as the history associated with both. This is a foundational part of the storytelling but it's also helpful to better understand how this technology relates to the world around us. For many, understanding AI/ML concepts can be mystifying, and unless you come from a tech or computer science background, the topics themselves can seem intimidating. Many will, at best, only acquire a rudimentary understanding of what this tech is and how it's come about.

We want to empower anyone interested in exploring this underlying tech that will shape so many products and internal systems in the future by making a deeper understanding more accessible. Already, it may seem like there's a bias – most of the folks who intimately understand ML and DL already come from a computer science background, whether it's through formal education or through boot camps and other technical training programs. That means that, for the most part, the folks who have had access to this knowledge and pursued study and entrepreneurship in this field have traditionally been predominantly white and predominantly male.

Apart from the demographics, the level of investment in these technologies, from an academic perspective, has gone up. Let's get into some of the numbers. Stanford University's AI index states that AI investment at the graduate level among the world's top universities has increased by 41.7%. That number jumps to 102.9% at the undergraduate level. An extra 48% of recipients of AI PhDs have left academia in the past decade in pursuit of the private sector's big bucks. Also, while only 14.2% of computer science PhDs were AI-related 10 years ago, that number is now above 23%. The United States, in particular, is holding onto the talent it educates and attracts. Foreign students who come to the United States to pursue an AI PhD stay at a rate of 81.8%.

The picture this paints is one of a world that's in great need of talent and skill in AI/ML. This high demand for the AI/ML skill set, particularly a demographically diverse AI skill set, is making it hard for people to stay in academia, since the private sector handsomely rewards those who don't. In the start-up circuits, many venture capitalists and investors are able to confidently solidify their investments when they know a company has somebody with an AI PhD on staff, whether or not their product needs this heavy expertise. Placing a premium on human resources with these sought-after skills is likely not going to go away anytime soon.

Another key thing to remember is the relationship between ML teams and product teams. How the models are chosen, built, tuned, and maintained for optimized performance is the work of data scientists and ML engineers. Using this knowledge of performance toward the optimization of the product experience itself is the work of PMs. If you're working in the field of AI product management, you're working incredibly closely with your data science and ML teams (we'll learn more about specific roles later in the book).

We'd like to also make a distinction about the folks you'll be working with as an AI PM. Depending on your organization, you're either working with data scientists and developers to deploy ML or you're working with ML engineers who can both train and upkeep the models as well as deploy them into production. We highly suggest maintaining strong relationships with any and all of these impacted teams, along with DevOps.

We dream of a world where people from many competencies and backgrounds come into the field of AI because diversity is urgently needed and the opportunity that's ahead of us is too great for the gatekeeping that's been going on to prevail. It's not just important that the builders of AI understand the underlying tech and what makes its application of it so powerful. It's equally important for the business stakeholders that harness the capabilities of this tech to also understand the options and capabilities that lie before them. At the end of the day, nothing is so complicated that it can't be easily explained.

ML versus DL – understanding the difference

In this section, we will explore the relationship between ML and DL and the way in which they bring their own sets of expectations, explanations, and elucidations to builders and users alike. Whether you work with products that incorporate ML models that have been around since the 50s or use cutting-edge models that have sprung into use recently, you'll want to understand the implications either way. Incorporating ML or DL into your product will have different repercussions. Most of the time, when you see an AI label on a product, it's built using ML or DL, so we want to make sure you come out of this chapter with a firm understanding of how these areas differ and what this difference will tangibly mean for your future products.

As a PM, you're going to need to build a lot of trust with your technical counterparts so that, together, you can build an amazing product that works as well as it can technically. We will go over some of the basics here, and will be elaborating on these concepts later on. Let's first take a look at how some of the key concepts that we'll be discussing are interlinked:

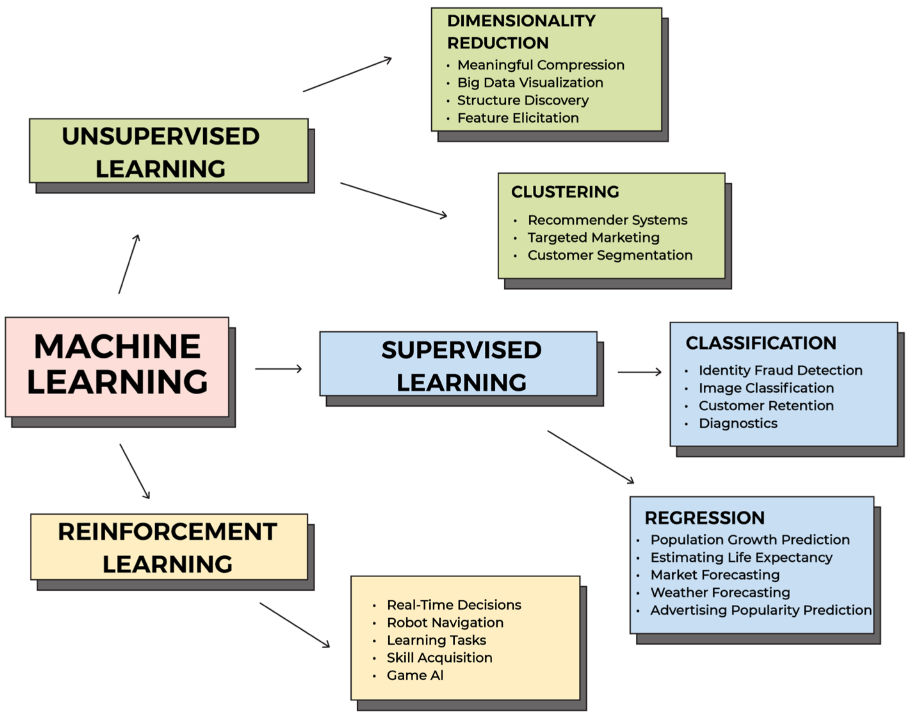

Figure 1.5: The subcategories of ML

ML

In its basic form, **ML** is made up of two essential components:

- The models used
- The training data it's learning from

These data are historical data points that effectively teach machines a baseline foundation from which to learn, and theoretically, every time you retrain the models with fresh data, the models improve. Retraining alone, however, doesn't guarantee better performance. It will improve performance if the fresh data is higher quality data, if it's more representative of real-world conditions, or if it's better labeled than the original data it trained on.

All ML models can be grouped into the following four major learning paradigms:

- Supervised learning
- Unsupervised learning
- Semi-supervised learning
- Reinforcement learning

These are the four major areas of ML, and each area is going to have its particular models and algorithms that are used in each specialization. The learning type has to do with whether or not the ML algorithm is learning from labeled or unlabeled (structured or unstructured) data, as well as the method you're using to reward the models you've used for good performance. These learning paradigms are relevant whether your product is using a DL model or not, so they're inclusive of all ML models. We will be covering the learning paradigms in more depth later in this chapter.

> Keep in mind that the "learning" is coming from improvements upon past mistakes. ML models are trying to learn from past performance to keep getting better with time. We use a few key metrics to keep track of how a model is learning and improving: accuracy, precision, and recall. We will go into these metrics further as the book goes on, but for now, you can keep these terms for reference:
>
> - **Accuracy**: This tells you how often your predictions are correct. It's like asking, "Out of all the guesses I made, how many times was I right?" For example, if you correctly identify both good and bad cases most of the time, you have high accuracy.
> - **Precision**: This measures how accurate your positive predictions are. Imagine you're identifying something specific, like picking out all the ripe apples in a basket. Precision answers, "Of the apples I said were ripe, how many actually were?" High precision means you're good at picking out the right ones without including too many wrong ones.
> - **Recall**: This shows how well you're finding all the positive cases. It's like asking, "Of all the ripe apples in the basket, how many did I actually find?" High recall means you're good at not missing any of the ones you want to find.

DL

DL is a subset of ML, but the terms are often used colloquially as almost separate expressions. The reason for this is DL is based on neural network algorithms and ML can be thought of as... the rest of the algorithms. DL refers to models that have neural networks and all other models (including language models and computer vision models) are referred to as ML. In the preceding section covering ML, we looked at the process of taking data, using it to train our models, and using that trained model to predict new future data points. Every time you use the model, you see how *off* it was from the correct answer by getting some understanding of the rate of error so you can iterate back and forth until you have a model that works well enough. Every time, you are creating a model based on data that has certain patterns or *features*.

This process is the same in DL, but one of the key differences of DL is the depth of the models – patterns or features in your data are largely picked up by the DL algorithm through what's referred to as **feature learning** or **feature engineering** through a hierarchical layered system. Here's a diagram showing how it works:

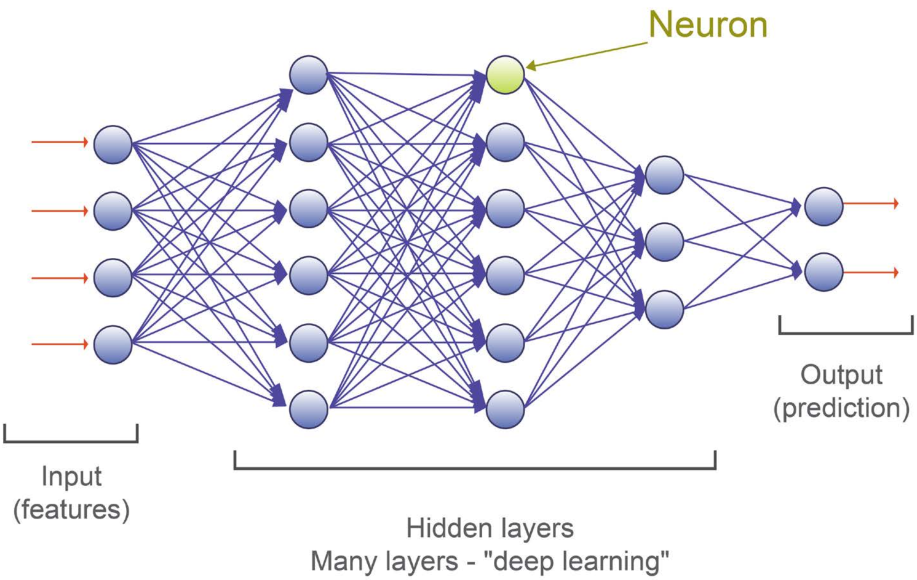

Figure 1.6: How ANNs work

We will go into detail about the various algorithms that are used in *Chapter 3* because they have a few nuances, but as you continue developing your understanding of the types of ML out there, you'll also start to group the various models that make up these major areas of AI (ML and DL). For marketing purposes, you will, for the most part, see terms such as *ML*, *DL/neural networks*, or just the general umbrella term of *AI* referenced where DL algorithms are used.

It's important to know the difference between what these terms mean in practice and at the model level and how they're communicated by non-technical stakeholders. As AI PMs, we are toeing the line between the two worlds: what engineering is building and what marketing is communicating. As mentioned previously, anytime you've heard the term **black-box model**, it's referring to a neural network model, which is DL. The reason for this is DL engineers often can't determine how their models are arriving at certain conclusions, creating an opaque view of what the model is doing. This opacity is double-sided, both for the engineers and technologists, as well as for the customers and users downstream who are experiencing the effects of these models without knowing how they make certain determinations.

For AI PMs, DL poses a concern for explainability because there's very little we can understand about how and why a model is arriving at conclusions. Depending on the context of your product, the importance of explainability could vary. One of the biggest objections your sales team might face from customers is that they're not able to purchase black-box model products; this came up for us quite a bit in some of our own work experiences. Another inherent challenge is these models essentially learn autonomously because they aren't waiting for their engineer to choose the features that are most relevant in the data for them; the neural networks themselves do the feature selection. They learn with very little input from an engineer. Think of the models as the *what* and the following section of learning paradigms as the *how*.

Learning paradigms in ML

In this section, we will cover the differences between supervised, unsupervised, semi-supervised, and reinforcement learning and how all these learning categories can be applied. Again, the learning type has to do with whether or not you're labeling the data and the method you're using to reward the models you've used for good performance. The ultimate objective is to understand what kind of learning model gets you the kind of performance and explainability you're going to need when considering whether or not to use it in your product.

Supervised learning

If humans are labeling the data (also known as structured data) and the machine is also looking to correctly label current or future data points, it's supervised learning. Because we humans know the answer the machines are trying to arrive at, we can see how off they are from finding the correct answer, and we continue this process of training the models and retraining them until we find a level of model performance that we're happy with.

Applications of supervised learning models include classification models that are looking to categorize data in the way spam filters do or regression models that look for relationships between variables in order to predict future events and find trends. Keep in mind that all models will only work to a certain point, which is why they require constant training and updating, and AI teams often use ensemble modeling or will try various models and choose the best-performing one. It won't be perfect either way, but with enough handholding, it will take you closer and closer to the truth.

The following is a list of common supervised learning models/algorithms you'll likely use in production for various products:

- **XGBoost**: Used for classification and regression tasks, this model is named after "extreme" gradient boosting, meaning it employs a variety of models that try to correct the errors from the previous. It works by minimizing the loss function of each model as optimally as possible, making it a popular choice for many supervised learning applications. It's commonly used for predicting customer churn or detecting fraud.
- **Naive Bayes classifier**: This algorithm *naively* considers every feature in your dataset as its own independent variable. So, it's essentially trying to find associations probabilistically without having any assumptions about the data. It's one of the simpler algorithms out there and its simplicity is what makes it so successful with classification. It's commonly used for binary values such as trying to decipher whether or not something is spam or understanding sentiment.
- **Support vector machine** (**SVM**): This algorithm is also largely used for classification problems and will essentially try to split your dataset into two classes so that you can use it to group your data and try to predict where future data points will land along these major splits. If you're not seeing compelling groups between the data, SVMs allow you to add more dimensions to be able to see groupings more easily. Commonly used for image or text classification.
- **Linear regression**: These models have been around since the 1950s and they're the simplest models we have for regression problems such as predicting future data points. They essentially use one or more variables in your dataset to predict your dependent variable. The *linear* part of this model is trying to find the best line to fit your data, and this line is what dictates how it predicts. Here, we once again see a relatively simple model also being heavily used because of how versatile and dependable it is. Commonly used for predicting prices or forecasting sales.
- **Logistic regression**: This model works a lot like linear regression in that you have independent and dependent variables, but it's not predicting a numerical value; it's predicting a future binary categorical state, for instance, whether or not someone might default on a loan in the future. Commonly used for diagnosing diseases or customer behavior.
- **Decision trees**: This algorithm works well with both predicting something categorical (colors or groups) as well as something numerical (age or income), so it's used for both kinds of ML problems, such as predicting a future state (categorical) or a future price (numerical), which has contributed to its popularity. Its comparison to a tree comes from the nodes and branches that effectively function like a flow chart. The model learns from the flow of past data to predict future values. Commonly used for automating loan approvals or segmenting customers based on historical behaviors.
- **Random forest**: This algorithm builds from the previous decision trees and is also used for both categorical and numerical problems. The way it works is it splits the data into different *random* samples, creates decision trees for each sample, and then takes an average or majority vote for its predictions (depending on whether you're using it for categorical or numerical predictions). It's hard to understand how a random forest comes to conclusions, so if interpretability isn't super high on the list of concerns, you can use it. Commonly used for feature selections in a dataset or evaluating creditworthiness for loan applications.

- **K-nearest neighbors** (**KNNs**): This algorithm exclusively works on categorical as well as numerical predictions, so it's looking for a future state and it offers results in groups. The number of data points in the group is set by the engineer/data scientist, and the way the model works is by grouping the data, determining characteristics the data shares with its neighbors, and giving its best guess based on those neighbors for future values. Commonly used for image recognition and recommendation systems.

> Keep in mind that supervised learning may not always be accessible for all AI/ML teams. You may have heard of "weak supervision," an approach where ML models are learning from data that has been labeled by other AI systems. This approach is increasingly accessible and useful in cases where high-quality labeled data is scarce or expensive to acquire.

Now that we've covered supervised learning, let's discuss unsupervised learning next.

Unsupervised learning

If the data is unlabeled (also known as unstructured data) and we're using machines to label the data and find patterns we don't yet know of, it's unsupervised. Effectively, we humans either know the right answer or we don't, and that's how we decipher which camp the ML algorithms belong to. As you might imagine, we take the results of unsupervised learning models with some hesitancy because it may be finding an organization that isn't actually helpful or accurate. Unsupervised learning models also require large amounts of data to train on because the results can be wildly inaccurate if they're trying to find patterns out of a small data sample. As they ingest more and more data, their performance will improve and become more refined over time, but once again, there is no "c*orrect"* answer but, rather, grades of correctness. For instance, over time, we may be able to create zones or clusters in which we're looking to find data. Or perhaps it's the proximity of some data being within certain boundaries that we're looking for.

Applications of unsupervised learning models include the following:

- **Clustering** models segment or group data into certain areas. In the case of clustering, you're grouping data that's similar somehow, the idea being that, over time, you're able to identify distinct clusters or groups within data based on similarities. These can be used for things such as looking for patterns in medical trials or drug discovery, for instance, because you're looking for connections and groups of data where there might not already be obvious answers.
- **K-means clustering** is a common unsupervised learning algorithm that will group data points together to better see patterns (or clusters), but it's looking for some optimal number of clusters as well. This is unsupervised learning, so the model is looking to find patterns that it can learn from because it's not given any information (or supervision) to go off from the engineer that's using it. Also, the number of clusters assigned is a hyperparameter and you will need to choose what number of clusters is optimal. Common uses for K-means include market segmentation (grouping customer groups and profiles) and document clustering (grouping news articles into topics or categories).

- **Dimensionality reduction** essentially removes the features in your dataset that contribute less to the performance you're looking for and will simplify your data so that your most important features will best improve your performance to separate real signals from the noise. In cases where you have high dimensional data, you might have too many features and you're looking for a way to maintain the variance or structure from your original data but through fewer features.
- **Principal component analysis** (**PCA**) is an unsupervised learning model that is often used for dimensionality reduction. Often, the largest problem with using unsupervised ML on very large datasets is there's actually too much uncorrelated data to find meaningful patterns. This is why PCA is used so often because it's a great way to reduce dimension without actually losing or discarding much information. This is especially useful for massive datasets such as finding patterns in genome sequencing or drug discovery trials. Common uses for PCA include data visualization (reducing high dimensional data into 2D or 3D landscapes) and reducing features in a dataset for improving performance and efficiency in other ML models.

> We always use the term **hyperparameters** when defining model optimizations because "parameters" refer to the boundaries within the training data that the model is using to make predictions. When it comes to adjustments to the model and how it functions, the term will always "be hyperparameter".

Next, let's jump into semi-supervised learning.

Semi-supervised learning

In a perfect world, we'd have massive labeled datasets with which to create optimal models that don't overfit. **Overfitting** is when you create and tune a model to the dataset you have but it fits a bit too well, which means it's optimized for that particular dataset and fails to generalize on unseen data. The issue of overfitting is a common problem in data science and applies to all the models we're discussing in this chapter. We live in an imperfect world and we can find ourselves in situations where we don't have enough labeled data or enough data at all. For instance, perhaps your company is launching a new product or service and there's little to no historical data to train models on user behavior, sales trends, or performance metrics. This is where semi-supervised learning comes in handy. We give some labeled datasets and also include a dataset that is unlabeled to essentially give the model nudges in the right direction as it tries to come up with its own semblance of finding patterns. It doesn't quite have the same level of absolute truth associated with supervised learning, but it does offer the models some helpful clues with which to organize its results so that it can find an easier path to the right answer.

For instance, let's say you're looking for a model that's trying to label pets in Google Photos. You might label a few of them and then see how the performance improves over time with the examples you don't label. You can use multiple models in semi-supervised learning. The process would be a lot like supervised learning, which learns with labeled datasets so that it knows exactly how off it is from being correct. The main difference between supervised learning and semi-supervised learning is that you're predicting a portion of the new unlabeled data and then, essentially, checking its accuracy against the labeled data. You're adding unlabeled new data points into the training set so that it's *training* on the data it's gotten correct.

Semi-supervised learning techniques are less common and require specialized training and expertise to do correctly, which is due to a number of factors. Availability of correctly labeled data, the complexity of implementation, the sensitivity to data quality issues and noise, and variability in performance are all reasons why teams might be resistant to semi-supervised learning. There are also fewer tools, libraries, and solutions for semi-supervised learning compared to supervised and unsupervised learning.

Finally, to wrap up this section, let's take a brief look at reinforcement learning.

Reinforcement learning

This area of ML effectively learns with trial and error, so it's learning from past behavior and adapting its approach to finding the best performance by itself. There's a sequence to reinforcement learning and it's really a system based on weights and rewards to reinforce correct results. Eventually, the model tries to optimize for these rewards and gets better with time. We see reinforcement learning used a lot with robotics, for instance, where robots are trained to understand how to operate and adjust to the parameters of the real world with all its unpredictability. For example, models like those that power ChatGPT are initially trained with supervised learning but are then refined through **reinforcement learning from human feedback** (**RLHF**).

In reinforcement learning, you're learning from interactions that are happening within a large state space. Since labels aren't available here, these models are learning from sequences of actions, states, and rewards. This state space is typically quite large, so it's hard to explore all possible states and actions; some version of human feedback being introduced into the loop helps with managing the potentially exponential scale that can come from that state space. Reinforcement learning models are instead looking for an optimal learning path based on the feedback they receive. The more feedback given in the training process, the quicker they will learn.

Here is a brief summary of all the learning paradigms:

- **Supervised learning:** The model is trained on a labeled dataset. Common tasks include classification (spam detection, image recognition) and regression (predicting housing prices).
- **Unsupervised learning:** The model is trained on an unlabeled dataset. Common tasks include clustering and dimensionality reduction.
- **Semi-supervised learning:** The model is trained with a hybrid dataset, typically a smaller subset of labeled data with a larger subset of unlabeled data. Common tasks are similar to supervised learning but are most often used where data is scarce or hard to obtain like transcribed audio.
- **Reinforcement learning:** Learning is done in the context of an environment where rewards and penalties are given based on the accuracy of outputs. Common tasks include game playing, robotics, and autonomous systems like self-driving cars.

Now that we've discussed and understood the different ML types, let's move on to an area of AI that's received a lot of attention in the last few years: LLMs and generative AI.

LLMs, NLP, GANs, and generative AI

Just as "AI" is an umbrella term, "generative AI" follows suit. As you might have inferred from its name, generative AI is an area of AI that's all about generating new content, whether that's text, an image, or even code. The ML models that power generative AI are creating outputs that closely resemble the training data they learn from.

When you think of generative AI, think primarily of advanced DL models; these can be grouped into three major categories:

- **Latent variable models**: These models try to decipher hidden factors (latent variables) from the data they are given to train on. From the visible data they receive, they try to understand the determinants that are hidden in that data. Examples of latent variable models include **variational autoencoders** (**VAEs**) and **energy-based models** (**EBMs**). Examples of tools that are made with these kinds of models include Artbreeder, NSynth (Google Magenta), AutoCAD's Generative Design, and DeepLog. Netflix's recommendation system is also partially supported by VAEs.
- **Adversarial models**: These models work more like a game in that there are two adversarial models "fighting" in tandem to arrive at the best result. One model generates data while the other evaluates what it's generating to ensure it's worthy enough to offer as an output. Examples of adversarial models include **generative adversarial networks** (**GANs**) and adversarial EBMs. Examples of tools that are made with these kinds of models include DeepArt, Jukebox (OpenAI), Adobe Photoshop (Content-Aware Fill), DALL-E, GANimation, GameGAN (NVIDIA), FaceApp, and Prisma.
- **Sequential models:** These models have gotten good at predicting outputs in a sequence and are able to use previous elements contextually. They've paved the way for long-form text generation because they are able to generate one step at a time by considering all that came before the output they're trying to create. Examples of sequential models include autoregressive models, Transformer models, and diffusion models. Examples of tools that are made with these models include GPT models that power ChatGPT (OpenAI), Smart Compose (Google), MuseNet (OpenAI), PixelCNN, and Google Translate.

One of the most famous generative AI models to reach the market, ChatGPT, which was trained by OpenAI, was only launched in November 2022 after almost a decade of work done by a massive open-source community of developers and scientists. Now that it's been launched, many of the tech giants are trying to compete in the market with similar products. Even though they are struggling with creating something truly competitive given ChatGPT's prevalence in the market, some competing products have emerged including Google's Bard and Gemini, Microsoft's Bing Chat, and Anthropic's Claude and Mistral. Think about that for a moment: these models are so complex and require so much compute power and so much data that they're a formidable challenge even for the kinds of tech companies that are wealthier than most countries' GDP.

What makes these models special is that they're DL models that are accessible to people with very little understanding of AI and deep tech. The reason for that is when they're released to market, they are powered by prompts. This means you have to actually ask it a question or give it a "prompt" to get a result. Because of the quantity of data, most of which is written text, images, and web page content, they're good at understanding both what someone is asking for and also what someone might want to read or see back. The process by which these DL models understand a question and generate a response is called **natural language processing** (**NLP**).

Chatbots or **conversational AIs** are computer programs designed to simulate human conversation. They use **NLP** to understand and respond to text or voice inputs from users. Think of them as virtual assistants that handle customer service tasks like answering questions, providing recommendations, or helping with customer support in different ways. . NLP is made up of two components:

- **Natural language generation** (**NLG**): AI interprets and makes sense of the text or speech it receives by recognizing the meaning behind words, identifying known entities, and understanding the intent behind a user's message.
- **Natural language understanding** (**NLU**): AI then crafts coherent and contextually appropriate responses, such as simple acknowledgments, more detailed instructions, or more in-depth information based on what it has understood using NLU.

The models that harness this NLP capability are called LLMs. Most LLMs use a learning process that combines unsupervised, supervised, and reinforcement learning. In the case of ChatGPT, the model is pretrained on a massive dataset with unstructured data.

For those interested in an example dataset used in the unsupervised pretraining step, feel free to explore one here: `https://pile.eleuther.ai/`.

The stages involved are as follows:

1. Its first stage is unsupervised learning.
2. The next stage of its training process is supervised fine-tuning, in which it undergoes a corrective phase where a human reviewer steps in to tell it when it gets something right.
3. In the final learning phase, it undergoes reinforcement learning so that a system of rewards is introduced into its feedback loop to further improve its performance.

We will be going into further nuances about the relevance and details when it comes to generative AI models and the DL models that power them in later chapters of this book. But for now, this was a concise summary of the basics when it comes to generative AI. Before we discuss the optimal flow of the ML process, let's see what it takes to succeed in AI.

Succeeding in AI – how well-managed AI companies do infrastructure right

It's indicative of the complexity of ML systems that many large technology companies that depend heavily on ML have dedicated teams and platforms that focus on building, training, deploying, and maintaining ML models. The following are a few examples of options you can take when building an ML/AI program:

- **OpenAI Platform from OpenAI**: OpenAI's AI management platform encompasses tools and systems for developing, deploying, and managing AI models. This includes the training infrastructure, deployment pipelines, and monitoring systems they use to support the creation and operations of models like GPT-4 and other AI technologies they have.
- **MLflow from Databricks**: MLflow is an open source platform developed by Databricks to help manage the complete ML life cycle for enterprises. It allows you to run experiences and work with any library, framework, or language. The main benefits are:
    - Experiment tracking allows you to see how your models are doing between experiments.
    - Model management helps you manage all versions of your model between teammates.
    - Model deployment gives you a quick view of deployment in the tool.
- **TensorFlow Extended (TFX) from Google**: This is Google's newest product built on TensorFlow and it's an end-to-end platform for deploying production-level ML pipelines. It gives you the following benefits:
    - Collaboration within and between teams
    - Robust capabilities for scalable, high-performance environments
- **Michelangelo from Uber**: Uber is a great example of a company creating their own ML management tool in-house for collaboration and deployment. Earlier, they were using disparate languages, models, and algorithms and had teams that were siloed. After they implemented Michelangelo, they were able to do the following:
    - Bring in varying skill sets and capabilities under one system.
    - Create, manage, predict, and deploy their data at scale using a reliable, recreatable, and standardized pipeline.
- **FBLearner Flow from Meta**: Meta also created its own system for managing its numerous AI projects. Since ML is such a foundational part of its product, Meta needed a platform that would allow the following:
    - Every ML algorithm that was implemented once to have the ability to be reusable by someone else at a later date.
    - Every engineer to have the ability to write a training pipeline that can be reused.
    - Make model training easy and automated.
    - Everybody to have the ability to search past projects and experiments easily.

Effectively, Meta created an easy-to-use knowledge base and workflow to centralize all their ML ops.

- **SageMaker from Amazon**: This is Amazon's product that allows you to build, train, and deploy your ML models and programs with their own collection of fully managed infrastructure tools and workflows. The idea of this product is:
    - It allows them to meet their customers where they are and offer low-code or no-code UIs, whether they employ ML engineers or business analysts.
    - The ability to use their infrastructure is also great if you're already using Amazon services for your cloud infrastructure so that you can take it a step further to automate and standardize your ML/AI program and operations at scale.
- **Bighead from Airbnb**: Airbnb created its own ML infrastructure in an effort to create standardization and centralization between their AI/ML organizations. They used a collection of tools such as Zipline, Redspot, and DeepThought to orchestrate their ML platform in an effort to do the same as Meta and Uber:
    - Mitigate errors and discrepancies.
    - Minimize repeatable work.

As we can see, there are multiple platforms that can be used to create, train, and deploy ML models. Now, let's move on to the optimal flow of setting up an AI system (whether for the product or internally).

The order – what is the optimal flow and where does every part of the process live?

Companies interested in *creating value* with AI/ML have a lot to gain compared to their more hesitant competitors. According to McKinsey Global Institute, "*Companies that fully absorb AI in their value-producing workflows by 2025 will dominate the 2030 world economy with +120% cash flow growth.*" The undertaking of embracing AI and productionizing it – whether in your product or for internal purposes – is complex, technical debt-heavy, and expensive. Once your models and use cases are chosen, making that happen in production becomes a difficult program to manage, and this is a process that companies in industries other than tech might struggle with as they start to take on the challenge of embracing AI.

Operationalizing the process, updating the models, keeping the data fresh and clean, and organizing experiments, as well as validating, testing, and the storage associated with it, are the complicated parts. In an effort to make this entire process more digestible, we're going to present this as a step-by-step process because there are varying layers of complexity, but the basic components will be the same. Once you have gotten through the easy bit and you've settled on the models and algorithms you feel are optimal for your use case, you can begin to refine your process for managing your AI system.

> Keep in mind that the following steps will be expanded on in future chapters at greater length and depth. But for now, the following sections are a good reference for what you can expect when setting up and maintaining an AI system.

Step 1 – Definition

While much of this book will be focused solely on AI as it pertains to product management, you can use this process for implementing any AI system, even if it's not specifically integrated into your product. As any PM knows, you don't start out with a solution without first addressing the problem you're trying to solve. No matter the purpose of your AI system, you're going to want to articulate the specific problem you want to use AI for and the goal you're trying to achieve with it. This is what allows you to attach that problem to the business impact its potential solution could contribute to. Aligning an AI system to business impact and how it will help achieve business goals is the first step because you're going to ask your organization to put in a significant investment for setting up an AI system.

Before they grant your request, they'll have a lot of questions. Figuring out how to answer those questions and how to establish success criteria will be your primary task. Scoping out the bulk of the work and establishing requirements and constraints will be part of your upfront investment in creating an AI system. Being able to identify limitations and requirements relating to infrastructure, data availability, and compute power will ensure you're preparing yourself adequately as an AI PM. It won't just be the leadership's approval you need but stakeholders' approval as well. In most cases, you'll need to consult with your key stakeholders to be able to offer a reasonable and circumspect answer to the problem you're trying to solve.

Step 2 – Data availability and centralization

Essentially, you'll need a central place to store the data that your AI/ML models and algorithms will be learning from. Depending on the databases you invest in or legacy systems you're using, you might have a need for an **extract**, **transform**, **load** (**ETL**) pipeline and data engineering to make the layers of data and metadata available for your productionized AI/ML models to ingest and offer insights from. Think of this as creating the pipeline needed to feed your AI/ML system. Determine the type and amount of data you'll need to train and test your AI models so that you can gather the right data from various sources in order to establish your centralized, integrated training dataset.

AI feeds on data, and if your system of delivering data is clunky or slow, you'll run into issues in production later. Choosing your preferred way of storing data is tricky in and of itself. You don't know how your tech stack will evolve as you scale, so choosing a cost-effective and reliable solution is a mission in and of itself, as is cleaning and processing the data you do decide to use. You'll also want to split your dataset into a training, validation, and testing set. Depending on the model you choose and the purpose, you may also choose to augment your dataset with generated data points or external data sources.

For example, as we started to add more and more customers at a cybersecurity company we were previously working for, we noticed the load time for certain customer-facing dashboards was lagging behind. Part of the issue was the number of customers, and their metadata was too large to support the pipelines we already had in place.

Step 3 – Choose and train the model

Now that you've got a handle on the problem you're trying to solve with your AI system and how you're going to gather and use data, you can now try to build the solution. In this case, that's the model type you'll work with. Your data scientist(s) will be able to select an ML model class that would work well for the nature of the problem you selected and can work on researching suitable models that are known to perform well for the specific problem, type of data, computational resources, and transparency your organization is comfortable with.

This is also where you'll establish performance metrics for your model and get a sense of how long it takes to train and make predictions. This will all depend on a number of factors and it allows for the experimentation needed to settle on a good choice. You'll want to make sure you're doing your due diligence to test variants of models and configurations to find the best fit. This is where hyperparameter tuning and cross-validation come in to make sure you're making a sound choice in choosing the model(s) that will power your AI system. Based on the needs of your product team or organization, you'll settle on some combination that balances performance with complexity and scale, as well as resources.

Once you've settled on a suitable model or combination of models, you're going to prepare the resources you need to get your environment ready for training that model. In this stage, you're fitting the model to the training data to understand patterns and relationships in your data, and optimizing the performance of your model by adjusting hyperparameters. You're using the validation dataset at regular intervals to make sure the model is generalizing well with the data you have and you're also maintaining versions of your model as you train and validate.

Step 4 – Feedback

At this stage, you're evaluating the performance metrics you started out with and how those metrics perform against the testing dataset. Throughout this process, you're ensuring that the model that you chose is performing well and is generally strong enough to perform well against the validation and training datasets so that it's generalized enough to work with real-world data in production. Remember that your model will always make errors; no model is ever perfect. This step is all about how to understand and learn from the errors your model is making to establish a baseline of performance to measure from when your model is being used live in production.

At this step, you'll want to make sure you're investigating the types of errors and where they come from to identify patterns and causes of failure so that you can determine how to improve the model before you deploy it into your production environment. At this stage, you're also going to try to minimize your errors as much as you can through hyperparameter tuning, feature engineering, model compression, performance optimization, and final testing on edge cases. All of this will be thoroughly documented by the data scientist working with your product organization.

Step 5 – Deployment

Depending on the size of your product organization, you'll likely have data scientists working in tandem with ML engineers who will take all the work that was done on the model side and translate that into a successful deployment into production. This is where the deployment environment will be prepared for the trained model to be integrated into an existing system or application that is already in production. Models are typically integrated into these systems through APIs or web services. This is also where logs that continuously track the model's performance and operational status are set up.

You won't be able to understand how a model is performing if you're not using tools to monitor performance metrics like response time, accuracy, and errors, as well as system health and performance. This is also where security measures are put in place to keep the model and the data it's using encrypted, secure, and accessible by authorized users.

Step 6 – Continuous maintenance

At this point, you have your models and algorithms and you've chosen a system for delivering data to them. Now, you're going to be in the flow of constantly maintaining this system. In DevOps, this is referred to as **continuous integration** (**CI**)/**continuous delivery** (**CD**). In the later chapters, we will cover the concept of **AI for IT operations** (**AIOps**) but, for now, the following is a list of the stages tailored for the continuous maintenance of AI pipelines. The following are the four major components of the continuous maintenance process:

- **CI**: Testing/validating code and components, along with data, data schemas, and models.
- **CD**: Code changes or updates to your model are passed on continuously so that once you've made changes, they are slated to appear in the testing environment before going to production without pauses.
- **CT**: We've mentioned the idea of continuous learning being important for ML, and **continuous training** productionizes this process so that as your data feeds are refreshed, your models are consistently training and learning from that new data.
- **CM**: We can't have ML/AI models continuously running without also using **continuous monitoring** to make sure something isn't going horribly wrong.

You can't responsibly manage an AI program if you aren't iterating your process constantly, or without diligently documenting each step of this process. Your models and hyperparameters will become stale. Your data will become stale, and when an iterative process like this stagnates, it will stop being effective. Your documentation itself will become stale. Performance is something you'll constantly be staying up to date on because the lack of performance will be self-evident, whether it is client-facing or not. With that said, things can also go wrong.

There are severe ethical reasons for staying focused on proper AI system management. Depending on the product and its use, lags in performance or frequency of the model refreshing can lead to people losing their jobs, not getting a competitive rate on a mortgage, or getting an unfair prison sentence. Major consequences can arise from downstream effects due to improper model maintenance. We recommend exploring the *Additional resources* section at the end of this chapter for more examples and information on how stagnant AI systems can wreak havoc on environments and people.

Storing and managing data

AI/ML products run on data. Where and how you store your data is a big consideration that impacts your AI/ML performance, and in this section, we will be going through some of the most popular storage vehicles for your data. Figuring out the optimal way to store, access, and train your data is a specialization in and of itself, but if you're in the business of AI product management, eventually, you're going to need to understand the basic building blocks of what makes your AI product work. In a few words, data does.

Because AI requires big data, this is going to be a significant strategic decision for your product and business. If you don't have a well-oiled machine, you're going to run into snags that will impair the performance of your models and, by extension, your product itself. Having a good grasp of the most cost-effective and performance-driven solution for your particular product, and finding the balance within these various facets, is going to help your success as a PM. Yes, you will depend on your technical executives for a lot of these decisions, but you'll be at the table helping make these decisions, so some familiarity is needed here. Depending on your organization's goals and budget, you'll be centralizing your data somehow between a data lake, a database, and a data warehouse, and you might even be considering a new option: a data lakehouse. Let's look at some of these different options to store data for AI/ML products.

Database

If you're just getting your feet wet, you're likely just storing your data in a relational database so that you can access it and query it easily. Databases are a great way to do this if you have a relatively simple setup. If your primary use of the database is querying to access data and use only a certain subset of your company's data for general trends, a relational database might be enough. With a relational database, there's a particular schema you're operating under if you want to combine this data with data that's in another database. However, you may run into problems aligning these schemas later. So if you're looking to combine various datasets from disparate areas of your business and you're looking to accomplish more advanced analytics, dashboards, or AI/ML functions, you'll need to read on.

Common challenges with databases are that they can create data silos, so data can get fragmented across different systems or areas of a business. Combining that data can be time-consuming and difficult. I've worked for companies that only used databases but they weren't able to make the most of their data with those limitations. It wasn't until I worked for a company that was newly transitioning to using a data warehouse that I truly saw the full extent of what setting one up can offer a company. If you plan on querying your data and finding trends, particularly if you're using historical data to predict future trends, you'll need to have a way to centralize all the data you're storing in operational databases. For that, you'll need a data warehouse.

Data warehouse

If you're looking to combine data into a location where you can centralize it somewhere and you've got lots of structured data coming in, you're more likely going to use a data warehouse. This is really the first step toward maturity because it will allow you to leverage insights and trends across your various business units quickly. If you're looking to leverage AI/ML in various ways rather than one specific specialized way, this will serve you well.

Let's say, for example, that an e-commerce company collects a wide range of data from its website, sales transactions, customer behavior, inventory, and marketing campaigns. In order for that company to adequately understand past, current, and future trends and be able to report on those trends, they would need to find a way to use the wealth of data they have. A data warehouse would allow them to extract data from various sources like their sales system, **customer relationship management** (**CRM**), inventory management system, and marketing tools. It would also allow them to clean and transform the data from all these disparate sources so that they're all in the same standard format and tagging system. Finally, this new standardized data would be loaded into the data warehouse where it would be properly organized and stored for analysts to work with it. The importance of a centralized repository can't be understated, particularly if you want to be using advanced analytics to be able to forecast trends, predict customer behavior, employ AI for your product or for your business in the future, or even just properly use a business intelligence tool or workbench.

Data warehouses do, however, require some upfront investment to create a plan and design your data structures. They also require a costly investment because they make data available for analysis on demand, so you're paying a premium for keeping that data readily available. Depending on how advanced your internal users are, you could opt for cheaper options, but this option would be optimal for organizations where most of your business users are looking for easily digestible ways to analyze data. Either way, a data warehouse will allow you to create dashboards for your internal users and stakeholder teams.

Data lake (and lakehouse)

If you're sitting on lots of raw, unstructured data, and you want to have a more cost-effective place to store it, you'd be looking at a data lake. Here, you can store unstructured, semi-structured, and structured data that can be easily accessed by your more tech-savvy internal users. For instance, data scientists and ML engineers would be able to work with this data because they would be creating their own data models to transform and analyze the data on the fly, but this isn't the case at most companies.

Keeping your data in a data lake would be cheap if you've got lots of data your business users don't need immediately, but you won't ever really be able to replace a warehouse or a database with one. It's more of a "nice to have" for companies that have massive stores of "big" data. In the previous section, we covered the benefits of a data warehouse. To put it simply, data warehouses are able to take raw data and refine it to the point of making it usable. The reason why a company wouldn't want to use a data lake as a replacement is because data lakes are optimized for large volumes of raw data, not transformed data. Also, if you need data on demand, data lakes will likely be too slow for you.

Many data lakes have a "schema-on-read" approach, so they apply structure to the data when it's observed rather than having a latent schema. Again, because the data is raw, no schema will be applied to the data. This is why it's great for storing but not so great for accessing the data. Even if you did want to access it, it's good practice to load it into a data warehouse for analysis. If you're sitting on a massive data lake of historical data you want to use in the future for analytics, you'll need to consider another way to store it to get those insights.

You might also come across the term **lakehouse**. There are many databases, data warehouses, and data lakes out there. However, the only lakehouse we're aware of has been popularized by a company called Databricks, which offers something like a data lake but with some of the capabilities you get with data warehouses, namely, the ability to showcase data, make it available and ingestible for non-technical internal users, and create dashboards with it. The biggest advantage here is that you're storing it and paying for the data to be stored upfront with the ability to access and manipulate it downstream.

Data pipelines

Regardless of the tech you use to maintain and store your data, you're still going to need to put up pipelines to make sure your data is moving, that your dashboards are refreshing as readily as your business requires, and that data is flowing the way it needs to. There are also multiple ways of processing and passing data. You might be doing it in batches (batch processing) for large amounts of data being moved at various intervals, or in real-time pipelines for getting data in real time as soon as it's generated. If you're looking to leverage predictive analytics, enable reporting, or have a system in place to move, process, and store data, a data pipeline will likely be enough. However, depending on what your data is doing and how much transformation is required, you'll likely be using both data pipelines and perhaps, more specifically, ETL pipelines:

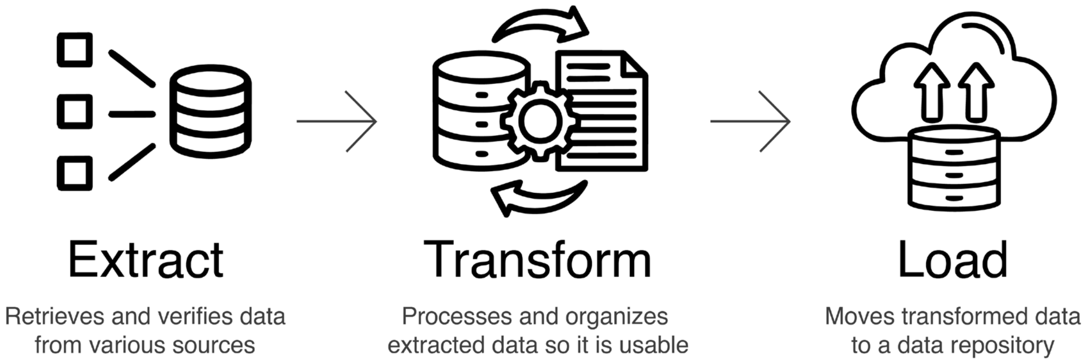

Figure 1.7: The ETL process (icons obtained from Flaticon.com)

ETL stands for **extract, transform, and load**, so your data engineers are going to be creating specific pipelines for more advanced systems such as centralizing all your data into one place, adding data or data enrichment, connecting your data with CRM tools, or even transforming the data and adding structure to it between systems. The reason for this is that it's a necessary step when using a data warehouse or database. If you're exclusively using a data lake, you'll have all the metadata you need to be able to analyze it and get your insights as you like. In most cases, if you're working with an AI/ML product, you're going to be working with a data engineer who will power the data flow needed to make your product a success because you're likely using a relational database as well as a data warehouse. The analytics required to enable AI/ML features will most likely need to be powered by a data engineer who will focus on the ETL pipeline.

Managing and maintaining this system will also be the work of your data engineer, and we encourage every PM to have a close relationship with the data engineers who support their products. ETL pipelines are generally updated in batches and not in real time. If you're using an ETL pipeline, for instance, to update historical daily information about how your customers are using your product to offer client-facing insights in your platform, it might be optimal to keep this batch updating twice daily. If you need insights to come in real time for a dashboard that's being used by your internal business users that rely on that data to make daily decisions; however, you likely will need to resort to a data pipeline that's updated continuously.

Now that we understand the different available options to store data and how to choose the right option for the business, let's discuss how to manage our projects.

Managing projects – IaaS

If you're looking to create an AI/ML system in your organization, you'll have to think about it as its own ecosystem that you'll need to constantly maintain. We will start to see managed services and **infrastructure-as-a-service** (**IaaS**) offerings coming out more and more as time goes on, particularly with the rise of AI being implemented internally for company optimizations but also for product optimizations. There has been a shift in the industry toward companies such as Determined AI and Google's AI Platform Pipeline tools to meet the needs of the market. At the heart of this need is the desire to ease some of the burden from companies left scratching their heads as they begin to take on the mammoth task of getting started with an AI system.

Just as DevOps teams became popular with at-scale software development, the result of decades of mistakes, we will see something similar with MLOps and AIOps. Developing a solution and putting it into operation are two different key areas that need to work together. This is doubly true for AI/ML systems. The trend now is on IaaS. According to Gartner, generative AI and application modernization are the reasons "IaaS is forecast to experience the highest end-user spending growth at 25.6%" (2024) and it's projected to grow 29.1% in 2025. This is an important concept to understand because companies just approaching AI often don't have an understanding of the cost, storage, compute power, and investment required to do AI properly, particularly for DL AI projects that require massive amounts of data to train on.

> **Artificial intelligence operations** (**AIOps**) focus on using ML and AI to automate and enhance IT operations and infrastructure for things like issue detection and predictive analytics to prevent outages, manage incidents, analyze root causes, or optimize system performance.
>
> **Machine learning operations** (**MLOps**) focus on optimizing and managing the life cycle of ML models from development and deployment to monitoring and maintenance. The goal is to ensure models are efficiently integrated into production environments.

At this point, most companies haven't been running AI/ML programs for decades and don't have dedicated teams. Tech companies like Meta, Amazon, Apple, Netflix, and Google are leading the cultural norms with managing AI/ML, but most companies that will need to embrace AI are not in tech and are largely unprepared for the technical debt that AI adoption will pose for their engineering teams to manage. Managing and integrating legacy systems alongside new AI technologies and capabilities, building structure around the AI/ML life cycle, bringing in enough technical skill to make it all happen in the first place, scoping compute resources and storage needs, managing cloud and on-premises services and infrastructure, and managing costs are all complexities an AI PM will need to think about and prepare for.

Shortcuts taken to get AI initiatives off the ground will require code refactoring or changing how your data is stored and managed, which is why strategizing and planning for AI adoption is so crucial. Even beyond the planning and costs, an AI PM will need to help the organization measure the success and failure of the AI initiatives it's planning, particularly if those AI initiatives are directly supporting their product. Setting up performance metrics, measuring the impact on the business, establishing a pattern of adoption and satisfaction, and linking that to **return on investment** (**ROI**) are all important facets of an AI PM's work. This is why so many of these IaaS services are popping up to help keep engineering teams nimble should they require changes in the future as well. The infrastructure needed to keep AI teams up and running is going to change as time goes on, and the advantage of using an IaaS provider is that you can run all your projects and only pay for the time your AI developers are actually using data to train models.

Deployment strategies – what do we do with these outputs?

Once you're happy with the models you've chosen (including their performance and error rate) and you've got a good level of infrastructure to support your product and chosen AI model's use case, you're ready to go to the last step of the process and deploy this code into production. Keeping up with a deployment strategy that works for your product and organization will be part of the continuous maintenance we outlined in the previous section. You'll need to think about things such as how often you'll need to retrain your models and refresh your training data to prevent model decay and data drift. You'll also need a system for continuously monitoring your model's performance so this process will be really specific to your product and business, particularly because these periods of retraining will require some downtime for your system.

> Over time, an ML model's performance can experience a gradual reduction in performance if there are changes to underlying patterns or relationships in the data that is being used to train them. This is known as **model decay**.
>
> When data used by an ML model changes or shifts from what the model was originally trained on, it can make a model's predictions less accurate because the patterns have changed. This is known as **data drift**.

Deployment is going to be a dynamic process because your models are trying to effectively make predictions of real-world data for the most part, so depending on what's going on in the world of your data, you might have to give deployment more or less of your attention. For instance, when we were working for an ML property-tech company, we were updating, retraining, and redeploying our models almost daily because we worked with real estate data that was experiencing a huge skew due to rapid changes in migration data and housing price data due to the pandemic. If those models were left unchecked and there weren't engineers and business leaders on both sides of this product, on the client's end and internally, we might not have caught some of the egregious liberties the models were making on behalf of under-representative data.

There are also a number of well-known deployment strategies you should be aware of, which we will discuss now.

Shadow deployment strategy

In this deployment strategy (often referred to as **shadow mode**), you're deploying a new model with new features along with a model's that already exists so that the new model that's deployed is only experienced as a *shadow* of the model that's currently in production. This also means that the new model is handling all the requests it's getting just as the existing model does but it's not showing the results of that model. This strategy allows you to see whether the shadow model is performing better on the same real-world data it's getting without interrupting the model that's actually live in production. Once it's confirmed that the new model is performing better and that it has no issues running, it will then become the predominant model fully deployed in production and the original model will be retired.

For example, if a streaming company wants to test a new recommendation system without affecting its current users, it might deploy a new algorithm to run in parallel with the existing one. This means that the new algorithm could process the same user data that the original one does but it wouldn't actually recommend anything. The original algorithm would continue to offer recommendations while the new one they're trying out would still process the same information. That way, the product team could see how it would process data and perform compared to the one they already know works. This means that they would be able to test out the new algorithm and its effectiveness without disrupting the user experience of an established customer base.

A/B testing model deployment strategy

With this strategy, we're actually seeing two slightly different models with different features to get a sense of how it's working in the live environment concurrently. The two models are set up at the same time and the performance is optimized to reward conversion. This is effectively like an experiment where you're looking at the results of one model over another and you're starting with some hypothesis or expectation of how one is performing better than another, and then you're testing that hypothesis to see whether you were right. The differences in your models do, however, have to be slight because if there's too much variety between the features of the two, you actually won't understand what's creating the most success for you.

For example, if an e-commerce company wants to try out a new pricing algorithm to see if it would have an impact on sales, it could split users into two groups. One group would see the old pricing model while the other would experience the new pricing algorithm. Since these two groups would reflect real pricing and purchasing decisions, the results could be compared between the two over a period of time. This would allow the company to monitor the site and compare sales performance, user engagement, and conversions between the two groups of customers to be able to understand the real impact of the new pricing algorithm.

Canary deployment strategy

Here, we see a more gradual approach to deployment where you actually create subsets of users that will then experience your new model deployment. Here, we're seeing the number of users that are subjected to your new model gradually increasing over time. This means that you can have a buffer time between groups of users to understand how they're reacting and interacting with this new model. Essentially, you're using varying groups of your own users as testers before you release to a new batch so you can catch bugs more gradually as well. It's a slow but rewarding process if you have the patience and courage.

For example, if a photo-sharing social media app wanted to test out a new photo editing feature, it might only deploy that new feature to a small test group. In this case, the "canary" group. They could monitor the performance, adoption, and feedback from this small group to identify any bigger issues the feature or capability might have before they roll it out to their entire user base. Being able to test a feature with a smaller group minimizes potential risks they might otherwise have with a wider rollout. After the success of that initial experiment, they can have a chance to make changes or improvements from the feedback before opening it to a wider audience.

There are more strategies to choose from but keep in mind that the selection of these strategies will depend on the following:

- The nature of your product
- The budget, which is what's most important to your customers and users
- Your metrics and performance monitoring
- Your technical capacity and knowledge
- The timeline

Beyond deployment, you're going to have to help your business understand how often they should be doing code refactoring and branching as well.

Example

Let's take an example of a food delivery service that was looking to launch a new feature for an app that improved delivery time predictions. They considered their options between canary, A/B testing, and shadow deployment approaches with the following pros and cons:

- **Shadow deployment**: Running the new delivery time feature along with the current delivery time feature. The new feature would gather data about customer interactions but the current delivery time feature would be the only one actually outputting delivery times visibly to users. Let's look at the pros and cons of this approach:
    - **Pros**: The product team can test the new feature with real conditions without it actually affecting users.
    - **Cons**: Because users won't actually interact with the new delivery time feature, the product team won't actually know how they'd be affected by it and if it would impact user engagement or user behavior.
- **A/B testing**: The product team would split users into two groups with the first group using the current delivery time feature and the second using the new delivery time feature. The pros and cons would be as follows:
    - **Pros**: The product team can compare how effective the current and the new feature would be at predicting delivery times, as well as how they impacts the user, their behaviors, and their engagement.
    - **Cons**: Testing the features needs to be carefully considered and designed in a way where the two groups can be equitably compared. The product team also needs to be careful not to accidentally overlap users.
- **Canary**: The product team could test out the new delivery time feature with only 5% of its users and monitor their reaction to it. The following are the pros and cons:
    - **Pros**: This allows the product team to monitor any adverse reactions to user behavior and engagement and fix any issues before more users experience them.
    - **Cons**: The small group of users might not be reflective of their entire customer base. Some user segments that aren't in the canary group might be more sensitive to the new feature.

In the end, the product team decided to go with the *A/B testing deployment* strategy because they didn't want to risk unknowns from their canary group and they also didn't want to push a feature without any users directly experiencing the change like they would in the shadow deployment strategy. They wanted to get clear, comparative measures of the impact of the new feature on user engagement, user behavior, and delivery accuracy all at once. Seeing how the rollout would impact users in real-world scenarios and having a direct comparison to the existing feature was paramount. Because the company had the infrastructure to handle A/B testing appropriately, as well as the means to analyze the results comprehensively, they went ahead with an A/B test of their delivery time prediction feature.

Now that we've discussed the different deployment strategies, let's take a bird's eye view and see where AI is taking us.

The promise of AI – where is AI taking us?

So, where is this era of AI implementation headed and what does it mean for all industries? At this point, we're looking at an industry of geopolitical influence, a technologically obvious decision that comes with a lot of responsibility, cost, and opportunity. As long as companies and PMs are aware of the risks, costs, and level of investment needed to properly care for an AI program, use it as a source of curiosity, and apply AI/ML to projects that create success early on and build from that knowledge, those that invest in AI will find themselves experiencing AI's promise. This promise is rooted in quantifying prediction and optimization. Here are a few examples:

- 35% of Amazon's sales come from their personalization and recommendation engine because of its effectiveness at suggesting products based on their users' behaviors, preferences, and past purchases. As this is 35% of their total revenue, this has a substantial impact on both ROI for their AI investment as well as revenue.
- Netflix's content recommendation engine personalizes recommendations and improves user engagement and retention to the tune of $1B annually. By reducing churn and increasing customer retention, the company is able to reduce the costs of acquiring new customers as well as lost revenue from churning customers.
- Shopify uses AI to detect and prevent fraud in their transactions, helping their entire community of sellers and merchants feel safe on their platform. These tools have helped the company save from chargebacks and losses and have significantly decreased fraud costs with over 20 million fraudulent authorizations stopped from their interventions.

When a third of your revenue is coming from an AI algorithm, there's virtually no argument. Whatever investment you make in AI/ML, make sure you're leveraging it to its maximum capacity by properly planning and strategizing, finding capable talent that's aware of the space and potential dangers, and choosing the right scalable infrastructure to limit your refactoring.

As long as your AI/ML projects are directly married to outcomes that impact cost savings or revenue, you'll likely experience success within your own career if you're overseeing these projects. The recommendation of starting small, applying it to a clear business goal, tracking that goal, and showing off its effectiveness is a smart strategy because this chapter details the many areas of maintaining an AI program, as well as potential areas where it might experience hurdles. Justifying the time, investment in headcount, and infrastructure expenses will be challenging if you're not able to communicate the strength and capabilities of AI to even your most hesitant executive.

This will also be important for your *technical resources* (data scientists, data engineers, and ML engineers) as well as for your *business stakeholders*. It's one thing to know more about the ML algorithms you'll be using or to get a few recommendations about how to best store your data, but you really won't have the intimacy and fluency needed to truly be an agent of change within your organization if you don't iterate with your own projects and grow your knowledge of what works best from there. We learn through iteration and we build confidence the more we complete a task successfully. This will be the case for you as a PM as well.

In the examples discussed earlier in this section, Shopify prevented future bottlenecks by predicting fraud, Netflix reduced churn while increasing loyalty, and Amazon grew its revenue through ML. When we think about the promise of AI and where it's taking us, these examples drive the idea that this is the home of the latest industrial revolution. It's not just something that will offer benefits to companies but to everyone all at once. The distribution of the benefits may not be completely equal because, ultimately, it's the companies that are investing in this tech and they will look to experience the highest return on this investment first, but the point stands that consumers, as well as businesses, will experience benefits from AI.

Summary

We've covered a lot in this chapter, but keep in mind that this chapter serves as an introduction to the many terms and areas we will cover throughout the book. A lot of the concepts presented here will be returned to in subsequent chapters for further discussion. It's almost impossible to overstate that the infrastructure AI/ML will need to be successful because so much of the performance is dependent on how we deliver data and how we manage deployments. We covered the basic definitions of ML and DL, the learning paradigms that both can employ, as well as generative AI. We also covered some of the basics of setting up and maintaining an AI pipeline and included a few examples of how other companies manage this kind of operation.

Building products that leverage AI/ML is an ambitious endeavor, and this first chapter was meant to provide enough of a foundation for the process of setting up an AI program overall so that we can build on the various aspects of that process in the following chapters without having to introduce too many new concepts so late in the book. If you're feeling overwhelmed, it only means you're grasping the scale necessary for building with AI. That's a great sign! In *Chapter 2*, we will get into the specifics of using and maintaining the ML models we briefly introduced earlier in this chapter.

Additional resources

For additional information, you can refer to the following resources:

- *Weapons of Math Destruction* by Cathy O'Neil: `https://www.amazon.com/Weapons-Math-Destruction-Increases-Inequality/dp/0553418815`
- *Invisible Women: Exposing Data Bias in a World Designed for Men* by Caroline Criado Perez: `https://www.amazon.com/Invisible-Women-Data-World-Designed/dp/1419735217/`
- *The Ethical Algorithm: The Science of Socially Aware Algorithm Design* by Michael Kearns and Aaron Roth: `https://www.amazon.com/Ethical-Algorithm-Science-Socially-Design/dp/0190948205/`
- *Artificial Unintelligence: How Computers Misunderstand the World* by Meredith Broussard: `https://www.amazon.com/Artificial-Unintelligence-Computers-Misunderstand-World/dp/026253701X/`
- *Algorithms of Oppression: How Search Engines Reinforce Racism* by Safiya Umoja Noble: `https://www.amazon.com/Algorithms-Oppression-Search-Engines-Reinforce/dp/1479837245/`

- *Race After Technology: Abolitionist Tools for the New Jim Code* by Ruha Benjamin: `https://www.amazon.com/Race-After-Technology-Abolitionist-Tools/dp/1509526404/`
- *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* by Shoshana Zuboff: `https://www.amazon.com/Age-Surveillance-Capitalism-Future-Frontier/dp/1541758005/`
- *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* by Virginia Eubanks: `https://www.amazon.com/Automating-Inequality-High-Tech-Profile-Police/dp/1250074312/`
- *Data Feminism* by Catherine D'Ignazio: `https://www.amazon.com/Feminism-Strong-Ideas-Catherine-DIgnazio/dp/0262044005/`

References

- *Machine Learning category on Canva*: `https://www.canva.dev/blog/engineering/tag/machine-learning/`
- *Model deployment strategies*: `https://neptune.ai/blog/model-deployment-strategies`
- *3 ways Duolingo improves education using AI*: `https://blog.duolingo.com/ai-improves-education/`
- *GGWP platform*: `https://www.crunchbase.com/organization/ggwp-65c2`
- *Overview of Hazy*: `https://www.cbinsights.com/company/anon-ai`
- *AI–50 America's most promising artificial intelligence companies*: `https://www.forbes.com/lists/ai50/`
- *Shein's AI-driven expansion and the growing challenge of sustainability*: `https://greyjournal.net/news/shein-ai-expansion-sustainability-challenges-2023/`
- *Product Led Growth, Wes Bush*
- *Mind the gap – It's not AI/ML unless it's in production: Data strategy series Part 4*: `https://www.credera.com/insights/mind-gap-not-ai-ml-unless-production-data-strategy-series-part-4`
- *Airbnb's end-to-end ML platform*: `https://medium.com/acing-ai/airbnbs-end-to-end-ml-platform-8f9cb8ba71d8`
- *Amazon SageMaker*: `https://aws.amazon.com/sagemaker/`
- *Introducing FBLearner Flow: Facebook's AI backbone*: `https://engineering.fb.com/2016/05/09/core-data/introducing-fblearner-flow-facebook-s-ai-backbone/`
- *From predictive to generative – How Michelangelo accelerates Uber's AI journey*: `https://www.uber.com/blog/from-predictive-to-generative-ai/`
- *TFX is an end-to-end platform for deploying production ML pipelines*: `https://www.tensorflow.org/tfx`
- *Managed MLflow*: `https://www.databricks.com/product/managed-mlflow`

- *Future-proof your data strategy with Lakehouse*: https://www.databricks.com/resources/ebook/the-data-lakehouse-platform-for-dummies
- *Computing machinery and intelligence*: https://phil415.pbworks.com/f/TuringComputing.pdf
- *Key requirements for an MLOps foundation*: https://cloud.google.com/blog/products/ai-machine-learning/key-requirements-for-an-mlops-foundation
- *How does TikTok use machine learning?*: https://dev.to/mage_ai/how-does-tiktok-use-machine-learning-5b7i
- *Sparks of artificial general intelligence: Early experiments with GPT-4*: https://arxiv.org/pdf/2303.12712
- *The world's tech giants, compared to the size of economies*: https://www.visualcapitalist.com/the-tech-giants-worth-compared-economies-countries/
- *2024 AI Index Report*: https://aiindex.stanford.edu/report/

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

https://packt.link/aipm

2

Model Development and Maintenance for AI Products

In this chapter, we will be exploring the nuances of model development, from linear regression to deep learning neural network models. We'll cover the variety of models that are available to use, as well as what's entailed for the maintenance of those models, from how they're developed and trained to how they're deployed and ultimately tested. This will be a basic overview to understand the end-to-end process of model maintenance that product managers can expect from the engineering and dev ops teams that support their products.

There's a lot involved with bringing any new product to market, and if you've been a product manager for a while, you're likely familiar with the **new product development** (**NPD**) process – or set of steps. As a precursor to the rest of the chapter, particularly for those who are unfamiliar with the NPD process, we're going to be summarizing each of the steps in the first section of this chapter. Overall, this chapter will cover the following topics:

- Understanding the stages of NPD
- Model types – from linear regression to neural networks
- Training – when is a model ready for market?
- Deployment – what happens after the workstation?
- Testing and troubleshooting
- Refreshing – the ethics of how often we update our models

Understanding the stages of NPD

In this section, we will be covering the various stages of the NPD cycle as it relates to the emergence of an AI/ML product.

> While not every NPD is a net new product, product managers often enhance products that are already in production. We will review the stages here with the assumption that it is a new product.

Through each stage, we'll cover the major foundational areas, from the ideation to the launch of an acceptable first version of a product. The steps are laid out incrementally:

1. In the **Discovery** stage, you brainstorm about the need you're looking to address in the market and why that need needs to be bolstered by AI.
2. In the **Define** stage, you bring in your product requirements for your product.
3. In the **Design** stage, you bring in the active visual and experiential elements of your end product.
4. In the **Implementation** stage, you build it out.
5. In the **Marketing** stage, you craft a message for your broader audience.
6. In the **Training** stage, you put your product to the test and make sure it's being used as intended.
7. Finally, in the **Launch** stage, you release your product to a broader audience for feedback.

This can be illustrated by the following framework:

Figure 2.1: The NPD stages

Let's get into these stages in more detail in the following sections.

Step 1 – Discovery

In this phase, you're ideating. You look to isolate the particular problem you're trying to solve, and in the context of a **machine learning** (**ML**) product, a crucial part of this first phase is understanding why you're trying to solve that particular problem with ML in the first place. To borrow a phrase from Simon Sinek's popular book *Find Your Why*, this is where you "find your why." This is the phase in which you contemplate the fundamentals of the problem at hand and look to isolate what is most urgent about the problem so that you can later address an unmet need or under-served customers.

This requires gathering qualitative and quantitative customer feedback about the particular issue they're facing that you're looking to address, as well as doing market and competitive analysis to understand what potential solutions are out there today. The biggest focus here is creativity – to brainstorm potential solutions that you can then analyze and further explore (or discard) later. Not every problem will be solved with AI/ML and it's important to develop discernment when you're working with emerging technologies and be mindful and critical about when to employ them. In some cases, there might not be enough high-quality or diverse data to make AI/ML a reasonable solution to the problem you're trying to solve. Or, very simply, rules and formulas might be enough to solve your problem rather than training ML algorithms. In many cases, ML might be too costly a solution for the gains it may bring.

Step 2 – Define

The second phase is all about defining your **minimal viable product** (**MVP**). You've taken all the feedback about the problem and potential solutions in the first step, but now you're actually building a plan from those ideas. You have to start somewhere, right? So, this step is all about screening your ideas from the *Discovery* stage to select the one that has the highest potential to solve your customers' biggest problem. This is where all those creative brainstorming sessions are put to the test and analyzed to best understand which of the ideas from phase one have legs. What you're looking for here is the minimum number of features you'd need to create a version of your product that will address the main problem areas – or assumptions of – for the customers you're looking to serve.

As far as your model goes, this is also where you define some metrics for model performance that will mark the minimum performance your model will need to reach in order to be a good, viable option for your customer. Remember, this is just for your MVP. The idea is that you first begin with your MVP and then you iterate through sprints or product development processes to incrementally make your product perform better or build in features your customers might prefer or need over time. Model performance will work the same way. As you partner with your customers, you will refine the product, the models, and the performance of those models together over time.

Step 3 – Design

In the first and second steps of this process, you identify the problem you want to solve and whom you want to solve it for. Once you understand this, you'll come up with ideas, and then define the minimum amount of work you'll need to take on the problem at hand. Now, in this third *Design* step, you build out that MVP and start to piece together what it might look like. This is the step that's most heavy on finding the solution. In this step, you're coming up with mockups for how folks might interact with your product, what the UI might look like, and how the product experience might unfold. For AI products, this is also where you start to identify which of the models mentioned in the *Model types* section will best serve your product.

This step is all about creating a roadmap of the UI/UX elements if they're applicable to your product. Even in cases where it isn't, in the case of APIs, there are still design elements you'll need to consider like endpoint structure, naming conventions, and response formats. It's where you will want to involve some of your customers in the solution and, for an AI product, where you'll set some performance benchmarks and goals for your model to hit. Building performance into the design process and managing these expectations with eventual users of your product is a great way to crystalize the concept and test it early on.

Step 4 – Implementation

The *Implementation* phase is where all the ideating and planning from the first three steps are put to the ultimate test. This is the phase in which you're actually working to materialize everything you just worked hard on strategizing. For all intents and purposes, this is essentially your first sprint, and as a product manager, you're effectively working as a project manager in this phase to make sure that what you end up with meets the needs you set out to address.

While all your main stakeholders will be involved in all steps at different levels of engagement, in this stage, your engineers, ML engineers, developers, UI/UX folks, and project managers will be able to apply their work to create the MVP and achieve the performance your customers and the leadership are expecting. What you should be left with is a version of your MVP that does what you said it would do, as you said it would. You know you've succeeded with this step when your MVP meets the planning criteria you detailed in the *Define* stage.

Step 5 – Marketing

Marketing happens in the background of all these steps because even part of *step 1* relates heavily to marketing. Understanding the language of your customer, their needs, and their pain points is a huge prerequisite for getting your messaging right. Marketing is the delivery and communication of your message to your wider market base, and the reason why it's *step 5* is that you want to have a working MVP before you craft the official message that will go out for your current and prospective customers to see.

With AI products, marketing undergoes specific scrutiny because the AI market is heavily competitive and companies are in a communication quagmire. There's still quite a lot of hype and misunderstanding around AI. As with any new technology, the value proposition isn't always clear to the customer upfront. If you communicate too much about your product and which models make it worthy of the AI stamp, you're giving away too much of the secret sauce. If you communicate too little about the actual tech that's giving it AI/ML capabilities, you're likely to face criticism that you're overselling your solutions' AI capabilities.

We can say with a lot of confidence that most companies err on the side of under-communicating when it comes to AI products. Ultimately, the key to a successful marketing strategy is being able to articulate how this new technology will make a user's life better – the value it will bring to those who choose to adopt it. This is the step in which you will need to agree with all your stakeholders on how you best want to communicate AI capabilities to the outside world.

Step 6 – Beta testing

The process of testing a product on users and documenting your product happens in this sixth phase so that you can create the justifications for the choices you've made for your MVP and your product overall. Part of testing your product on your users is also managing expectations for how they are to interpret the performance of your product. Releasing a nearly complete product to a selected group of users (beta testers) is a way to gather feedback on its performance, usability, and potential issues in the real world post-launch. This part will be especially important with AI/ML products because they often optimize, rank, classify, recommend, or predict future values, and it will be especially important to help your customers understand when they can trust or question certain results.

This process is intuitive for the most part because when it comes to AI/ML, we don't know how far off we are from the predictions or optimizations until a future point in time. Part of the training that must happen, then, is to manage expectations with your customers about what margins of error are healthy for them to expect. This step is all about informing others about your product and how they can best interact with it. The goal of this step is to identify bugs or user experience issues early, before launch.

Step 7 – Launch

In this final step, we launch the product into the market officially. So far, you've spoken to your internal stakeholders and teams, received customer feedback, and maybe one or two customers have partnered with you to help create your offering and bring it to market. Maybe you've had a soft launch or gotten other beta testers/users to help you as well, but ultimately, the final step is your official hard launch. A big part of this final step is actually scaling back to your original definitions for performance and customer success. Is this final version of your product hitting the metrics you originally set with your customers? Is the performance of the product what everyone expected? Are you actively seeking to define future achievable goals? You will have gotten some feedback and likely sent out surveys in the *Marketing* and *Beta testing* phases, so you'll be able to link feedback you're getting post-launch to these earlier samples to gauge how your product is being received by your customers and users.

Now that we've covered the process that's commonly followed in NPD, we can move on to the models that are commonly employed in that development cycle. In the following section, we will review the most popular ML model types that are commonly used in production, as well as some of the characteristics those models share.

Model types – from linear regression to neural networks

In the previous chapter, we looked at a few model types that you'll likely encounter, use, and implement in various types of products for different purposes. To jog your memory, here's a list of the ML models/algorithms you'll likely use in production for various products:

- **Naive Bayes classifier**: This algorithm "*naively*" considers every feature in your dataset as its own independent variable, so it's essentially trying to find associations probabilistically without holding any assumptions about the data. It's one of the simpler algorithms out there and its simplicity is actually what makes it so successful with classification. It's commonly used for binary values, such as trying to decipher whether something is spam or not.
- **Support Vector Machine** (**SVM**): This algorithm is also largely used for classification problems and will essentially try to split your dataset into two classes so that you can use it to group your data and try to predict where future data points will land along these major splits. If you don't see compelling groups within the data, SVMs allow you to add more dimensions to be able to see groupings more easily.
- **Linear regression**: These models have been around since the 1950s and they're the simplest models we have for regression problems, such as predicting future data points. They essentially use one or more variables in your dataset to predict your dependent variable. The "linear" part of this model tries to find the best line to fit your data. This line, determined by minimizing the sum of the squared differences between observed and predicted values, is used to make predictions. Here, we once again see a relatively simple model heavily used because of how versatile and dependable it is.

- **Logistic regression**: This model works a lot like linear regression in that you have independent and dependent variables, but it doesn't predict a numerical value – it predicts a future binary categorical state, such as whether or not someone might default on a loan in the future, for instance.
- **Decision trees**: This algorithm works well for both categorical and numerical predictions, so it's used for both kinds of ML problems, such as predicting a future state or a future price. Decision trees are used often for both kinds of problems, which has contributed to its popularity. Its comparison to a tree comes from the nodes and branches, which effectively function like a flow chart. The model learns from the flow of past data to predict future values.
- **Random forest**: This algorithm builds from the previous decision trees and is also used for both categorical and numerical problems. The way it works is it splits the data into different "random" samples, creates decision trees for each sample, and then takes an average or majority vote for its predictions (depending on whether you're using it for categorical or numerical predictions). It's hard to understand how random forest comes to conclusions, so if interpretability isn't super high on the list of concerns, you can use it.
- **K-Nearest Neighbors** (**KNNs**): This algorithm exclusively works on categorical and numerical predictions, so it looks for a future state and offers results in groups. The number of data points in the group is set by the engineer/data scientist and the way the model works is by grouping the data, determining characteristics that data shares with its neighbors, and making the best guess for future values based on those neighbors.
- **K-means clustering**: This algorithm will group data points to see patterns (or clusters) better, but it looks for an optimal number of clusters as well. This is unsupervised learning, so the model looks to find patterns that it can learn from because it's not given any information (or supervision) to go off of from the engineer who's using it. Also, the number of clusters assigned is a hyperparameter, and you will need to choose what number of clusters is optimal.
- **Principal component analysis** (**PCA**): Often, the largest problem with using unsupervised ML on very large datasets is there's actually too much uncorrelated data to find meaningful patterns. This is why PCA is used so often, because it's a great way to reduce dimensions without actually losing or discarding information. This is especially useful for massive datasets, such as finding patterns in genome sequencing or drug discovery trials.
- **Neural networks**: Deep learning models are lumped under the term neural networks for the most part because they all mimic the way the human brain processes information through layers of nodes and their edges. There are several neural network types with their own particulars, but for now, it suffices to say that neural networks are what make up the models used in what we call **deep learning**. Deep learning is responsible for most current forms of generative AI that have been popularized in recent years. We briefly touched on generative AI in *Chapter 1*, and we will go deeper into the various models used for generative AI in the following chapter. In the meantime, here is a list of free deep learning resources where PMs can test models with their own data:
    - Hugging Face Model Hub: `https://huggingface.co/models`
    - TensorFlow Hub: `https://tfhub.dev/`

- PyTorch Hub: `https://pytorch.org/hub/`
- Google Colab: `https://colab.research.google.com/`
- Kaggle: `https://www.kaggle.com/`
- TorchVision Model Zoo: `https://pytorch.org/vision/stable/models.html`

If you see that a product is labeled as an AI/ML product, it likely uses some form or combination of these aforementioned types of models. We will be going over these models in later chapters of this book, but for now, this is a good introduction to the model types you'll most often come across where ML and AI are referenced. Now that we've introduced the models, let's go into how those models are trained and made ready for use in production.

OKRs

No matter what domain, vertical, or peer group your AI product is in, you're going to need to establish some way of communicating the success of your product through a combination of value (business) metrics, **key performance indicators** (**KPIs**), and **objectives and key results** (**OKRs**), along with a number of technical metrics that might be required when you're communicating about the efficacy and success of your product to a technical audience. As with anything, if we can't establish a baseline and see how we've grown from that baseline, we won't know whether our performance is improving (and if it is, by how much) unless we track it.

In the following section, we will be looking at the various types of metrics we will start to collect on our products' efficacy. For AI products, deciding on which metrics you will track, how you will talk about them, and what kinds of audiences you'll tailor certain metrics for will be an important part of your product strategy, as well as your marketing.

Objectives and key results

When you're defining success for your AI product, you'll want to set some OKRs from a technical and business level so you can track how your AI product is building toward the performance you want to see. OKRs are used heavily in product management to track progress toward higher-level business goals. They're a popular goal-setting framework that's standard for most PM roles (and beyond) and they help map team or organizational objectives to measurable outcomes.

OKRs should be transparent and should come with ambitious but achievable objectives. Having too many of them could dilute the focus of your team. You'll want to choose two or three objectives to start with, and you'll want a group of three to five key results for each objective. When working with OKRs, you should focus on the outcomes or "key results" as opposed to the tasks it took to get there. The purpose of this is to show which direct results are impacting the greater overarching business goals your product might be tasked with, not to prove you've got a to-do list.

Here is a sample OKR:

- **Objective:** Increase customer satisfaction by enhancing product quality
- **Key Result 1:** Improve the product's **net promoter score** (**NPS**) from 60 to 70 by the end of Q4.
- **Key Result 2:** Reduce customer support tickets by 30% by the end of Q4.

- **Key Result 3**: Reduce product downtime to less than 1% through improved QA and testing initiatives by the end of Q4.

You'll want to set these yearly and quarterly and reference them often because these will be the highest focus of your product as you build through each release. The establishment of new OKRs will arise in your product planning sessions and will be informed by the features you have planned in your roadmap, your customers, and customer success and marketing feedback, as well as the ever-evolving goals from leadership in order to best align your product's objectives with the feedback loops you have in place to make sure you're building something that's valuable to your domain and market.

Metrics and KPIs

KPIs are used more generally and also tend to grow over time as new indicators are discovered as important. Some of the most common KPIs have to do with measuring customer or employee satisfaction, measuring time or accuracy, calculating the cost efficiency or return on investment of purchasing something, or agreeing on a metric that aligns with the company's goals somehow. There are also categories of KPIs ranging from strategic to operational to functional, which means they're either affecting the entire company as a whole or are high-level metrics, relating to a specific time frame operationally across the entire company, or relating to specific departments.

If you've been working in the business world for some time, you've likely heard of KPIs or value metrics in a number of contexts, but in this section, we're going to specifically cover some common KPIs that are likely to help you when you're deciding on how to best track and communicate your AI product's success.

In *Chapter 1*, we discussed regression, classification and clustering, and deep learning ML models, as well as NLP. In this section, we will be expanding on those model categories to review the metrics and KPIs that are used to signal acceptable model thresholds for success. Whether you're running an AI program agnostically in your business or running an AI program as part of your AI product feature stack, the following KPIs will help with maintaining the health and progress of your AI infrastructure and communicate the success of your AI activities, a common part of **AI operations** (**AIOps**). Now we'll look at a few metrics that will help communicate the success of your AI activities, which is a common part of AIOps:

- **Classification metrics**: These metrics evaluate how well a model assigns labels or categories to data points to help assess the balance between correct predictions and errors. Some examples are:
    - **Accuracy**: This is the ratio of correctly predicted instances out of the total instances.
    - **Precision**: This is the ratio of true positives against all true positives and false positives combined. It indicates how often the model is correct at positive predictions.
    - **Specificity**: This is the ratio of true negatives against all true negatives and false positives combined. It indicates how often the model is correct at negative predictions.
    - **Recall**: Sometimes referred to as sensitivity, this is the ratio of true positives to the sum of true positives and false negatives. It measures the ability of a model to find all the relevant cases.

    - **F1 score**: This is the average of precision and recall. It balances both metrics into a single score.

- **Regression metrics**: These metrics measure the accuracy of a model in predicting continuous values and gauge how closely predictions match actual values. Some regression metrics are:

    - **Mean absolute error (MAE)**: This tracks the average rate of error. It measures the average magnitude of errors in a set of predictions without considering the direction of the error. Basically, how far off were the predictions from the actual outcomes. This metric is especially useful in scenarios where outliers are not as important.
    - **Mean squared error (MSE)**: This tracks the average of the squares of the errors. It gives a larger penalty to larger errors, which means that the higher the score, the further your predictions are from the actual values.
    - **Root mean square error (RMSE)**: The root of the average rate of error for regression models.
    - **R-squared**: This is a statistical measure that represents how good your model is at capturing the relationship between inputs and outputs.

- **Clustering metrics** : These metrics assess how well a model can group data points into clusters based on the similarities they share, as well as how clusters are separated from each other. The following are a few examples:

    - **Silhouette score**: This measures how similar an object is to its own cluster compared to other clusters. The higher the score, the better data points are clustered.
    - **Bavies-Bouldin index**: This measures the average similarity between each cluster and its most similar cluster. The lower the score, the better clustering you have between groups.
    - **Homogeneity**: This assesses whether clusters are effective in grouping similar items together. The higher the score, the more similar items in a cluster are.
    - **Completeness**: This assesses whether relevant items are captured in the same cluster. The higher the score, the more relevant items are included in the cluster.

- **NLP metrics**: There metrics evaluate the quality of models that process and generate text, measure the similarity of outputs to reference texts, or assess how natural and coherent text generation is. They can be of the following types:

    - **BLEU score**: This measures the similarity between a generated text and reference texts. It gives a quantitative way to evaluate the quality of the generated text and how closely it resembles human-generated text.
    - **Recall-Oriented Understudy for Gisting Evaluation (ROUGE) score**: This measures how good a model is at extracting (sentiment, for instance) and summarizing. It gives a way to assess how much important content is captured when text is being generated.
    - **Perplexity**: This is used for evaluating the predictive performance of language models. The lower the perplexity, the better the model is at predicting what the user wants.

- **IT operations and maintenance metrics:** These are metrics that are tracked in the maintenance of AI systems themselves:
    - **Mean time to detect (MTTD):** This calculates the average time it takes for your product to identify a potential issue. You'll want to demonstrate that this is minimizing over time.
    - **Mean time to acknowledge (MTTA):** This calculates the average time to acknowledge the problem and identify who will resolve it. You'll want to demonstrate that this is minimizing over time.
    - **Mean time to resolve/repair (MTTR):** This calculates the average time it takes to actually address the problem. You'll want to demonstrate that this is minimizing over time.
    - **Mean time between failures (MTBF):** This calculates the average time between failures to give a sense of how long the AI program is working optimally. You'll want to demonstrate that this is growing over time.
    - **Ticket to incident ratio:** Recognition that one incident may create multiple tickets, so this metric seeks to minimize the number of tickets or logs that are created by customers when one issue is experienced. You'll want to demonstrate that this is minimizing over time.
    - **Service availability:** The uptime where your AI program is running optimally without issues. You'll want to demonstrate that this is growing over time.
    - **Automated versus manual resolution:** This is essentially labeling which response was manual or automated so that your ML program can optimize for strategies in the future and learn from past remediations. You'll want to demonstrate that this is growing over time.
    - **User reporting versus automatic detection:** This is the ratio between how many customers reported an issue versus how often your own product detected the issue. You'll want to demonstrate that this is minimizing over time.
- **Technical metrics:** The following list includes some of the common technical KPIs that are used to communicate the accuracy of ML models:
    - **Classification accuracy (precision or specificity, recall or sensitivity, and F1 score):** There are a number of formulas that try to derive how often your models are correct, whether by the rate of true and false positives, the rate of true and false negatives, or some combination of the two (an F1 score).
    - **Root mean square error (RMSE):** The root of the average rate of error for regression models.
    - **Mean absolute error (MAE):** The average rate of error.
    - **Other metrics:** There are a number of other metrics you might want to keep track of. This may include metrics that pertain to the volume or amount of data sources or availability, perhaps you'll want to quantify or qualify the strength or robustness of your AI/ML organization somehow or have metrics related to your product marketing outputs or the availability of models.

Adopt metrics beyond just your users to understand how certain teams are using your products internally. Defining business goals is a highly selective and customized activity that should, first and foremost, serve the highest ideals of your business: to best serve your customers and market. We will discuss OKRs, metrics, and KPIs often throughout the rest of the chapters in this book – feel free to refer to this section for a refresher.

Training – when is a model ready for market?

In this section, we will explore the standard process for gathering data to train a model and tune hyperparameters optimally to achieve a certain level of performance and optimization. In the Implementation phase (*step 4* of the NPD process), we're looking for a level of performance that would be considered optimal based on the Define phase (*step 2* of the NPD process) before we move to the next phase of Marketing and crafting our message for what success looks like when using our product. A lot must happen in the Implementation phase before we can do that. Some of the key considerations are as follows:

- Data accessibility is the most important factor when it comes to AI/ML products. At first, you might have to start with third-party data, which you'll have to purchase, or public data that's freely available or easily scraped. This is why you'll likely want or need to partner with a few potential customers. Partnering with customers you can trust to stick with you and help you build a product that can be successful with real-world data is crucial to ending up with a product that's ready for market. The last thing you want is to create a product based on pristine third-party datasets or free ones that then becomes overfitted to real-world data and performs poorly with data coming from your real customers that it's never seen before.
- Having a wide variety of data is important here, so in addition to making sure it's real-world data, you also need to make sure that your data is representative of many types of users. Unless your product caters to very specific user demographics, you're going to want to have a model trained on data that's as varied as possible for good model performance as well as good usability ethics. There will be more on that in the final section.
- The next key concept to keep in mind with regard to training ML models is minimizing the loss function. While training data is key, your loss function is going to determine how off from the mark your model is performing. The process of training is exactly that: using data and adjusting your models to optimize for how correct it is at predicting an output. The more incorrect it is, the higher your loss. The more correct it is, the more you've minimized your loss function. The more your machine learns (and practices) the better its chances of good performance.
- Iterative hyperparameter tuning will also be hugely important as you continuously retrain your models for performance. One of the tools you have at your disposal, apart from changing/improving your training data, is adjusting the hyperparameters of your model. Note that not all models have hyperparameters to tune but most do. Models like linear regression models do have coefficients that can change, they are not at the discretion of the engineer. In contrast, deep learning models have the most hyperparameters and this is a big part of their training process.

- The performance metrics and benchmarks in the Define phase (*step 2* of the NPD) will inform how your ML engineers will go about tuning their hyperparameters. Most of the time, we don't yet know what the optimal model architecture for a certain use case is. We want to explore how a model functions with various datasets and start somewhere so that we can see which hyperparameters give us superior performance.
- Examples of what hyperparameters do include the degree of features that should be used in a linear model, the maximum depth that should be allowed for a decision tree model, how many trees should be included in a random forest model, or how many neurons or layers should be included for a neural network layer. In all these cases, we're looking at the external settings of the model itself and all these settings are worthy of scrutiny based on the model performance they produce. Having competent AI/ML engineers who are comfortable with navigating these shifts in performance will be important in creating a product that's set up for success.

We want to go into some applied examples of models and their comparisons to give product managers out there who are unfamiliar with AI/ML performance benchmarks a sense of how you can go about evaluating whether one model is better than another. The following are a few examples of performance metrics that your ML engineers will look at as they evaluate whether or not they're using optimal models. Note that not using optimal models could come with significant engineering and financial costs from the need to correct mistakes, including time and computational resources to redevelop and retrain your models.

You'll notice some of the names are familiar from our previous list of model types:

> These comparisons were done on a personal project, which was a model we had created to predict the price of Ether, a form of cryptocurrency. If you'd like to see the entire project outlined, you can do so here: `https://medium.com/analytics-vidhya/predicting-ether-prices-model-selection-for-machine-learning-8a50321f51a3`.

- The first model we wanted to use was an **ordinary least squares** (**OLS**) regression model because this is the most straightforward of the linear regression models that we wanted to select to give us a good baseline before we approached other model types.

  The results of the OLS regression model are as follows:

```
The number of observations in training set is 723
The number of observations in test set is 181
R-squared of the model in the training set is: 0.8985831338240027
-----Test set statistics-----
R-squared of the model in the test set is: 0.8896944272555466
Mean absolute error of the prediction is: 42.39871596133024
Mean squared error of the prediction is: 7572.23276225187
Root mean squared error of the prediction is: 87.0185771100164
Mean absolute percentage error of the prediction is: 90.68081323695678
```

Figure 2.2 – OLS regression model results

In *Chapter 1*, we discussed the notion of performance metrics for ML models and how to track them. There are a number of metrics that are automatically generated when you train a model. In the example above, we see what the full list of available metrics looks like when you run a model. For our comparison, we will be focusing on the **R-squared of the model in the test set** line in *Figure 2.2* to get the rate of error that's comparable between models. The **R-squared** metric is also referred to as the "coefficient of determination" and the reason why we use this particular metric so often in regression models is that it best assesses how far the data lies from the fitted regression line that the regression model creates. With the preceding OLS regression model, we see an R-squared of **0.889** for the test set using an 80/20 split of the training data. We used 80% of the data for training and the remaining 20% of the data for testing.

- The next model we tested was a random forest to compare results with a tree-based model. One of our hyperparameters for this random forest example was setting our cross-validation to `10` so that it would run through the training 10 times and produce an average of those 10 iterations as a final score. That average was an R-squared of 0.963, higher than our OLS model!

  The results of the random forest model are as follows:

```
cross_val_score(randomforest, X_test, Y_test, cv=10)
```

```
array([0.9491968 , 0.94922887, 0.97426398, 0.96202586, 0.97348678,
       0.99491192, 0.9764517 , 0.96363981, 0.96975411, 0.98030483])
```

```
import statistics

data = [0.96906062, 0.94844658, 0.94470685, 0.97056179, 0.97284841,
        0.98021631, 0.98151656, 0.95956996, 0.95165316, 0.94865387]

x = statistics.mean(data)
print(x)
```

```
0.962723411
```

Figure 2.3 – Random forest model results

- Finally, the last comparison was with our KNN model, which produced a score of 0.994. The hyperparameter we chose in this model was 6, which means we are looking for a group of 6 neighbors for each grouping. This KNN model gives us our best performance because we're ideally looking for the closest we can get to a perfect score of 1. However, we must keep this in mind with a caveat: although you are looking to get as close as you can to 1, the closer you get to 1, the more suspicious you should be of your model. For instance, the results of the KNN model are as follows:

```
from sklearn import neighbors
from sklearn import neighbors
from numba import jit
import numpy
import matplotlib.pyplot as pyplot
import seaborn
from sklearn.datasets import make_regression
from sklearn.model_selection import train_test_split
from sklearn.neighbors import KNeighborsRegressor

# Build our model.
knn = neighbors.KNeighborsRegressor(n_neighbors=6)
knn.fit(X_train, Y_train)
knn.score(X_test, Y_test)
```

```
0.994602173372774
```

Figure 2.4 – KNN model results

Though it may seem counterintuitive, getting this high a score likely means that our model is not working well at all, or that it's working especially well on the training data but won't perform as well on new datasets. While it may seem paradoxical, though the model is trying to get as close as it can to 1, getting too close is quite suspicious. That's because we always expect a model will be imperfect – there will always be some loss. When models perform exceedingly well with training data and get high scores, it could just mean that the model was calibrated to that data sample and that it won't perform as well with a new data sample.

This phenomenon is called **overfitting** and it's a big topic of conversation in data science and ML circles. The reason for this is that, fundamentally, all models are flawed and are not to be trusted until you've done your due diligence in selecting the best model. This game of choosing the right model, training it, and releasing it into the wild must be done under intense supervision. This is especially true if you're charging for a product or service and attempting to win the confidence of customers who will be vouching for you and your products someday. If you're an AI/ML product manager, you should look for good performance that gets better and better incrementally with time, and you should be highly suspicious of excellent model performance from the get-go. I've had an experience where model performance during training was taken for granted and it wasn't until we had already sold a contract to a client company that we realized the model performed terribly when applied to the client's real-world data. As a result, we had to go back to the drawing board and retrain a new model to get the performance we were looking for before deploying that model into our client's workflows.

A quick note on neural networks: while training generative AI models will be a bit different considering the subject matter and purpose of your model, it will follow a similar process. You're still going to put a premium on a clean and diverse data sample, you're still going to be thoughtful about which neural network will work best for the performance you want, and you're still going to need to account for (and optimize on) your loss function to the best of your ability. This process will continue through various loops of training and validating until you feel confident enough that your generative AI model will be able to generate new outputs based on the training examples you've given it. Your goal of tweaking hyperparameters for performance, minimizing loss where you can, and amassing enough data to set your model up for success remains the same as it does for other ML models.

Once you have comprehensive, representative data that you're training your models on, and you've trained those models enough times and adjusted those models accordingly to get the performance you're seeking (and promising to customers), you're ready to move forward!

Now that we've gone over some of the major aspects of model training, we can move on to what deployment looks like. Keep in mind that the entire process of ideating your product, choosing the right model to employ in your product, and gauging the performance of that model based on your training efforts is a collaborative effort. That collaboration doesn't end when you've trained your models; it intensifies. This is because you're now tasked with how exactly to integrate those models into the infrastructure of your product for your customers. Let's get into that in the following section.

Deployment – what happens after training?

In *Chapter 1*, we discussed deployment strategies that can be used as you manage your AI/ML products in production. In this section, we'd like you to understand the avenues available from a DevOps perspective, where you will ultimately use and deploy the models in production outside of the training workstation or training environment itself. Perhaps you're using something such as GitLab to manage the branches of your code repository for various applications of AI/ML in your product and experimenting there. However, once you are ready to make changes or update your models after retraining, you'll push the new models into production regularly. This means you need a pipeline that can support this kind of experimentation, retraining, and deployment regularly. This section will primarily focus on the considerations after we place a finished ML model into production (a live environment) where it will be accessed by end users.

How you manage these future deployments will vary widely depending on whether your AI/ML product offering is **business-to-business** (**B2B**) or **business-to-consumer** (**B2C**):

- If you're managing a B2C product, you'll likely make changes in phases and you'll likely use the deployment strategies outlined in *Chapter 1*, to manage how your updated product is received and when certain groups of users will see the new updated models. This is just the nature of a B2C product: it's one product going out to thousands, if not millions, of individual consumers, and your one product will mean many different things to individual users.
- If your product is a B2B product, then you manage expectations often at the customer level. One customer might have a different experience with your AI/ML product than another. The models you use could also very well change from one customer to another because the data you're using to train your models will be different from one customer to another.

Another thing to keep in mind is how you're going to handle discussions about your models and the collective training data you have among all your customers. With some products, you might not face much discussion about whether you use all your data to train your models. Some companies, however, are very particular about how their data is accessed and used. They might be okay with giving you historical data to train your model with as long as that data isn't being used to help the performance of other customers in their peer group, for example.

On the other hand, some customers might expect you to train your models on all the data you have to give your models the best shot at having as comprehensive a dataset as they possibly can. Remember that, as the strength of the models currently stands, the general rule is that the more data you have, the more examples you're able to give your models. This means that the more examples you have, theoretically, the stronger performance you should have across the board. Managing expectations with your customers and their threshold for data sharing is an important part of the deployment cycle because it's going to inform how often you update and how you deploy responsibly. While more data may make your model perform better overall, certain customers may be uncomfortable with their data being used in this way.

In one of my previous roles, we worked with a smaller number of **real estate investment trusts** (**REITs**). One day, I was on a customer call where we were discussing the idea of someday soon training our models on all our customer data combined, rather than having separate trained models for each customer. My point of contact on the customer side said that if we ever did that without their knowledge, they wouldn't renew with us the following year. The idea of their data being combined with their competitors' data in the service of our product offering made them so uncomfortable that they were willing to threaten our working relationship over it.

Managing expectations won't just be important on the client side of course. You'll likely have different teams that manage different areas of the deployment process. Perhaps your data scientists create and develop the models and train them, another team validates that work and the training data as well, and a third team of engineers deploys the models into the production environment. You might also have a team of ML engineers who specialize in different areas of this entire process.

Once you are ready to deploy your models, your ML engineers will analyze the deployment environment for the following reasons:

- To choose the best way to access the model (most often through an API or some UI/platform that's currently being used by your end user)
- To get a sense of how often it will be called
- To determine how many GPUs/CPUs and how much memory it will need to run
- To figure out how it will be continuously fed data

We'll leave the solutioning up to your onsite experts, but this is an important point for the AI/ML product manager to keep in mind: the time/money/effort/resources that will be required to keep your AI/ML algorithms running for your product will be a huge consideration when you choose the models and strategize how you'll deploy them in production.

The final part of deployment is training the end users on how to use the model and its results. Interpretability is important for any AI/ML project to succeed, but in the context of a product that's used and relied upon by end users, whether they are B2B (enterprise) or B2C customers, you will need to account for how to communicate through potentially confusing moments. Training your customers through in-app prompts or your customer success teams will allow your end users to learn how to activate your AI/ML features, access the data they need from these features, and interpret the output it gives them in a way that continuously reinforces your product's value – this is all part of managing your deployment.

Testing and troubleshooting

In *Chapter 1*, we discussed the idea of **continuous maintenance**, which included:

- Continuous integration
- Continuous delivery
- Continuous training
- Continuous monitoring

This section will build on that and expand on how to test and troubleshoot issues related to ML products on an ongoing basis so that your product is set up for success. Once you've done your first deployment, we jump right into the continuous training and continuous maintenance portion of the continuous maintenance process we discussed in *Chapter 1*.

Remember, managing the performance of your models post-deployment is crucial and it will be a highly iterative, never-ending process of model maintenance. As is the case with traditional software development, you will continue to test, troubleshoot, and fix bugs for your AI/ML products as well. The only difference is that you will also screen for lags in performance and bugs related to your model.

Continuously monitoring your model makes sure that it's always working properly and that the outputs it generates are effective. The last thing you want is for your product to be spewing out wildly inaccurate recommendations or predictions. Imagine that your model operated incorrectly and it took your customer weeks or months to notice that this had serious negative consequences downstream. They would question your integrity as a company because they trusted you to maintain and keep up the platform that they rely on for their own workflows. As a result, they might cancel their contract with you, pull their data out of your database, or give you negative reviews and negative referrals to other prospective customers.

Even when all the aspects of your model work properly, you will still need to track the continuous performance of your model and its outputs. The performance metrics for success we looked at in the *Training* section previously are the same metrics you'll create a log of that you routinely monitor to make sure model performance isn't lagging. In addition to statistical performance metrics, you'll also want to keep track of your accuracy, recall, and precision rates. This entire process of monitoring your model's performance should be automated so that you're alerted when certain metrics go over a certain threshold, in the form of a flag of some sort, so that you do not always have to manually check. For a larger breakdown of metrics, feel free to jump to the next section.

We don't just monitor the models themselves but we also continuously maintain the supporting code and documentation as well. This is notoriously a last priority for most companies that ultimately rely on the historical knowledge of the few developers that have been there the longest. Get it all documented and make it a practice of doing so regularly. You might find that there isn't enough training material or that the resources that currently exist just aren't adequate to explain what the product does. You might also find that the data feed that your model uses for training has issues with updates or wasn't properly connected in the first place. Perhaps it's an issue on your end users' side and they might not be accessing the AI/ML features of your product properly. Any number of these issues can happen routinely, which is why having teams devoted to the successful execution of your AI/ML product is crucial to its success.

Every model is going to have some form of degradation or **drift** over time. For example, if new data comes in that the model is training on that's not been cleaned in the same way that the training data was, your model's performance is going to suffer from a lack of uniformity in the data. **Data hygiene** is generally an important consideration when evaluating performance because it can wreak havoc and these kinds of changes might be hard to pinpoint.

Over months and years, if you see changes to how data is being reported and formatted, or if there are new fields or categories of data being added that weren't present when the models were first being trained, you're going to see the variance in your results. Data also can morph over time if your market changes or if the demographics of your users change. If major events impact the entirety of your dataset, this will adversely impact your model results because the baseline you built as your foundation will have been rendered unreliable because the majority of the training data might not apply to the new or current situation.

Outside of the training data, there is one last important area of drift, and that's what's often referred to as concept drift – or changes in your customer's expectations of what a correct prediction might be. For example, in some contexts, such as models predicting patient outcomes, they might face the concept of drift as new treatments become available or as patient populations change over time. Change is the only constant and the external environment is full of unpredictability. Any changes coming from outside factors could contribute to various types of concept drift, requiring us to go back to the drawing board, tweak our models, and redeploy them to address a changing world.

Continuous monitoring and testing are a big reason why many companies will use an enterprise data science platform to keep track of their deployments. We highly recommend this if you have the budget for it. Even if you don't, there are free or cheap open source tools out there to help you run experiments and maintain version control, like MLflow, DVC, Weights & Biases, Comet, and Neptune.ai. If you're working with many customers, as well as internal and external applications of your AI/ML models, you'll likely have many "reuse" use cases for your models. You'll benefit from the project tracking these platforms offer if you're managing at scale. For PMs looking to make a choice of platform, look for those that have the following features:

- Robust version control for models
- Easy tracking of changes and rollbacks
- Integration with your existing tools used by your data science team
- Detailed logs of each experiment you're tracking

In this section, we covered some of the most important considerations when testing and troubleshooting the use of your models in production, and the importance of regular monitoring for maintaining a level of oversight, not only to keep on top of the technical performance and robustness of your models but also to remain ethical. In the following section, we will focus more on the ethical considerations when building products with AI/ML components to build responsibly and harness some industry best practices.

Ethical retraining – the ethics of how often we update our models

When we think about the amazing power we have as humans, the complex brain operations we employ for things such as weighing up different choices or deciding whether or not we can trust someone, we may find it hard or impossible to believe that we could ever use machines to do even a fraction of what our minds can do. Most of us make choices, selections, and judgments without fully understanding the mechanism that powers those experiences. However, when it comes to ML, with the exception of neural networks, we can understand the underlying mechanisms that power certain determinations and classifications. We love the idea that ML can mirror our own ability to come to conclusions and that we can employ our critical thinking skills to make sure that process is as free from bias as possible.

The power of AI/ML allows us to automate repetitive, boring, uninspiring actions. We'd rather have content moderators, for instance, be replaced with algorithms so that humans don't have to suffer through flagging disturbing content on the internet on a daily basis. However, ML models, for all their wonderful abilities, aren't able to reason the way we can. Automated structures that are biased or that degrade over time have the power to cause a lot of harm when they're deployed in a way that directly impacts humans and when that deployment isn't closely and regularly monitored for performance. The harm that can cause at scale, across all live deployments of AI/ML, is what keeps ethicists and futurists up at night.

However, part of the danger with AI/ML is in the automation process itself. The types of drift we went over in the prior section impact how models derive meaning from the training data they learn from. Generative AI models in particular might be especially sensitive to this kind of stagnation because their purposes are so generalized. Since new data is constantly coming out with each passing day, the temptation to keep up with training them is very high. But so are the costs. Typically, LLMs are pre-trained during certain periods and those pre-trained models are then packaged and sold to companies, which then do their own training with their own specific data samples to create a more specialized model.

It is worth noting, however, that though ChatGPT has been out for some time, its last knowledge update as of March 2024 was actually January 2022. Though the English language itself may not have changed significantly since January 2022, the world as we know it has. People are increasingly using LLMs like ChatGPT for learning about the world around them and they're interpreting the outputs from these models as fact. Let's take, for instance, an AI PM using ChatGPT and taking its outputs as fact. It could impact their role and responsibilities by contributing to decision making bias, leading to decisions being made on inaccurate information. Decisions based on flawed information could lead to misguided product features, misaligned goals, or ineffective strategies.

The ethics with regard to generative AI models, particularly LLMs, need to account for the increased dependency on these models. I recently met with a recruiter who was telling me about schools like Southern New Hampshire University laying off members of their staff because they were using tools like ChatGPT to generate course content. While we can say that they likely (hopefully) have human editors who check the content of their courses for accuracy, if they feel confident enough to displace workers with LLMs, that doesn't diminish the importance of the stewards of LLMs to ethically maintain their models optimally considering the increased dependency on them.

The current state of accountability

Even when performance and maintenance appear normal, that doesn't mean that the models aren't taking liberties, resulting in real-world harm for the end user or for human beings that could be impacted downstream from the end user, whether or not they actually interact with the models themselves. A common example of this is the pervasive and unnecessary use of facial recognition software.

In February 2022, President Biden signed two pieces of legislation into law that expanded on AI accountability in the US: Artificial Intelligence for the Military Act of 2021 and the AICT Act of 2021. Will Griffin from *Fortune* magazine writes "*While this legislation falls far short of the calls for regulation consistent with the European Union model and desired by many in the A.I. ethics community, it plants the seeds of a thoughtful and inevitable A.I. ethics regulatory regime*." It's important to remember that AI ethics and regulations vary depending on where you live. In the US, we still lag behind European standards both in terms of legislation that's put in place to rein in AI misconduct and in terms of how we enforce existing laws.

On October 30th, 2023, President Biden signed an executive order on "Safe, Secure and Trustworthy Artificial Intelligence" in an effort to minimize the risks of AI. The new standards directed by this executive order require developers working on AI systems to transparently share test results and "critical information" with the US government and embolden the **National Institute of Standards and Technology** (**NIST**) to set testing and safety standards before AI systems are released, restrict AI-engineered biological materials, and limit AI-generated content with watermarks. Other notable aspects of the executive order revolve around protecting the privacy of American citizens' data, battling algorithmic bias and discrimination, consumer protections, mitigating worker harms brought on by AI, and the promotion of AI innovation and competition.

While executive orders are nice to have, their effects can vary and it's unclear how they'll be enforced moving forward. They're a step in the right direction, and the executive order on October 30th was lengthy and circumspect, but that's precisely what leaves them open to interpretation. Without legislation and hard regulations that are exact in definition, they can be rendered almost meaningless. Executive orders are also vulnerable to subsequent presidents coming in and limiting their language or revoking them altogether if they don't align with the current administration's priorities.

AI is still considered a wild west legislatively speaking, and we will likely see strides being made toward further defining the scope for how AI can interact with us as we see more and more use cases for AI products expand during this decade. AI PMs need to be aware of how legislation is changing at a local and global level to better prepare for potential changes to how AI products are conceptualized, built, and managed. Recently, the US made strides toward publishing a blueprint for an AI Bill of Rights that covers the following areas:

- Safe and effective systems
- Algorithmic discrimination protections
- Data privacy notices and explanation
- Human alternatives, consideration, and fallback

For now, we will use the European standards for framing how AI/ML product managers should think about their products because, even without deliberate laws that enforce AI ethics, entrepreneurs and technologists still face risks, such as losing customers, receiving bad press, or being taken to court, as a result of their algorithmic choices.

The European Commission outlines the following four key areas as ethical principles:

- **Respect for human autonomy**: "AI systems should not unjustifiably subordinate, coerce, deceive, manipulate, condition or herd humans. Instead, they should be designed to augment, complement and empower human cognitive, social and cultural skills. The allocation of functions between humans and AI systems should follow human-centric design principles and leave meaningful opportunity for human choice."
- **Prevention of harm**: "AI systems should neither cause nor exacerbate harm or otherwise adversely affect human beings. This entails the protection of human dignity as well as mental and physical integrity."
- **Fairness**: "While we acknowledge that there are many different interpretations of fairness, we believe that fairness has both a substantive and a procedural dimension. The substantive dimension implies a commitment to: ensuring equal and just distribution of both benefits and costs, and ensuring that individuals and groups are free from unfair bias, discrimination and stigmatisation."
- **Explicability**: "This means that processes need to be transparent, the capabilities and purpose of AI systems openly communicated, and decisions – to the extent possible – explainable to those directly and indirectly affected. Without such information, a decision cannot be duly contested. An explanation as to why a model has generated a particular output or decision (and what combination of input factors contributed to that) is not always possible. These cases are referred to as 'black box' algorithms and require special attention."

For more details, you can refer to *A framework for assessing AI ethics with applications to cybersecurity* by Danilo Bruschi and Nicla Diomede at `https://doi.org/10.1007/s43681-022-00162-8`.

On February 2nd 2024, EU member state representatives voted unanimously to approve the EU AI Act, making it the first comprehensive legal framework for AI anywhere in the world. Its impacts won't just be specific to EU member countries but to countless citizens all across the world. The AI Act will first group AI systems into various categories:

- Minimal risk
- Limited risk
- High risk
- Unacceptable risk

This is a concept significantly lacking in any of today's guardrails. Depending on the level or category of your AI system, responsibilities would include things like risk assessments, technical documentation and record keeping, transparency and disclosure, and compliance in accordance with the frameworks it will include. This is also a piece of legislation with teeth, with significant financial fines for companies/individuals in violation according to the AI category of their respective products.

For more details, you can refer to `https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai`.

Implementing ethical standards in your organization

Many companies might be tempted to create an AI ethics role within their companies and make it that person's problem or make them a scapegoat if and when they fail to meet certain standards, but this is a lazy and unethical way of managing the ethics around your AI programs if that's all you choose to do. A better way would be to train and empower all of the people who are involved in building your AI/ML products to be aware of the surrounding ethics and potential harm that could be caused to the customers or third parties interacting with your product.

While we must recognize the importance of understanding that recurring model updates are vital to maintaining good ethics with regard to ML and AI, as we've discussed previously in this chapter, it's also important to look into how your product can affect groups of people downstream who don't even use your product.

We don't exist in a vacuum. As we saw in the previous sections of this chapter, many factors at play already work against algorithms used in AI/ML products, which you have to keep track of even to stay on top of the natural chaos created by the constant input and output of data. This natural tendency that models have toward various types of drift is what demands a focus on ethics. According to a recent episode from TechTarget's *Today I Learned* podcast, FICO, the credit reporting and analytics vendor, conducted a survey of AI users and it showed that 67% of respondents do not monitor their models for accuracy or drift, which is pretty mind-blowing. These were AI users who were directly responsible for building and maintaining AI systems, which shows that the problems that come with unethical AI/data practices are the norm.

Ethical AI practices should be applied throughout every step we've outlined in this in-depth chapter on model development and maintenance. If we build AI/ML products that we are sure don't cause harm, both directly as part of our products' integrity and indirectly as part of our products' model maintenance, we can confidently market and promote our products without fear of retribution or punishment from the market that we want to serve. Every entrepreneur and technologist will have their own relationship with ethical business practices but, eventually, if you are a champion, promoter, or leader of a product that has come to market that harms others, you will be asked to explain what measures were put in place to inform your customers of the potential risks.

Summary

In this chapter, we covered the NPD cycle and a review of the common AI/ML model types. We also covered an overview of how to train, deploy, and troubleshoot the models that are chosen, giving us a reasonable foundation of what to expect when working with models in production. We also touched on some of the most important ethical practices, coming from some of the most rigorous standards that exist, when building products with AI/ML components.

If you're interested in expanding further on building ethical AI, we've provided some handy links in the following section for additional study. Keep in mind that we're at a critical juncture with regard to AI/ML ethics. We're building this ship as we're sailing it, and as AI/ML products continue to enter the zeitgeist, we will see additional measures put in place to reign in the potential harm caused by improper AI deployments through the diligent work of lawmakers and activists. We're not there yet, but with each new development, we get closer and closer to building a world that doesn't just embrace the promise of AI but limits the issues that AI poses as well.

So far, we've discussed the requirements of maintaining ML models and familiarizing ourselves with the process of building products with AI/ML components. These first two chapters are meant to serve as an introductory foundation in which we split deep learning from the broader umbrella term of ML so that we can get deeper into the concepts we've brought up so far in subsequent chapters. In *Chapter 3*, we'll focus on deep learning neural networks.

Additional resources

Reading about and familiarizing ourselves with AI ethics is important for everyone because AI is becoming increasingly impossible to avoid in our day-to-day lives. Additionally, if you actively work in the field of AI/ML as a data scientist, developer, engineer, product manager, or leader, it's doubly important that you're aware of the potential risks AI poses and how to build AI responsibly.

For further reading on ethical AI principles, we recommend the following reputable publications:

- *Blueprint for an AI bill of rights*: `https://www.whitehouse.gov/ostp/ai-bill-of-rights/`
- *DOD adopts ethical principles for artificial intelligence*: `https://www.defense.gov/News/Releases/release/article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/`

- *National artificial intelligence research and development strategic plan 2023 update*: https://www.whitehouse.gov/wp-content/uploads/2023/05/National-Artificial-Intelligence-Research-and-Development-Strategic-Plan-2023-Update.pdf
- *Fact sheet: President Biden issues executive order on safe, secure, and trustworthy artificial intelligence*: https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/
- *Algorithmic Justice League*: https://www.ajl.org/library/research
- AItruth.org's *12 Tenets of Trust*: https://www.aitruth.org/aitrustpledge
- Intel.gov's *Principles of AI ethics for the intelligence community*: https://www.intelligence.gov/principles-of-artificial-intelligence-ethics-for-the-intelligence-community
- European Commission's *Ethics guidelines for trustworthy AI*: https://www.aepd.es/sites/default/files/2019-12/ai-ethics-guidelines.pdf
- *UNESCO's Ethics of artificial intelligence*: https://en.unesco.org/artificial-intelligence/ethics
- "*Today I Learned*" podcast on ethical AI, with insights from Scott Zoldi, the Chief Analytics Officer at FICO: https://www.techtarget.com/searchcio/podcast/How-machine-learning-model-management-plays-into-AI-ethics

References

- *Find Your Why*, Simon Sinek: https://simonsinek.com/books/find-your-why/
- *High-level expert group on artificial intelligence set up by the European Commision*: https://www.aepd.es/sites/default/files/2019-12/ai-ethics-guidelines.pdf
- *Framing TRUST in artificial intelligence (AI) ethics communication: Analysis of AI ethics guiding principles through the lens of framing theory*: https://www.proquest.com/docview/2721197134
- *America must win the race for A.I. ethics*: https://fortune.com/2022/02/15/america-must-win-the-race-for-a-i-ethics-tech-artificial-intelligence-politics-biden-dod-will-griffin/
- *S.1776 – Artificial Intelligence for the Military Act of 2021*: https://www.congress.gov/bill/117th-congress/senate-bill/1776/text?q=%7B%22search%22%3A%5B%22s1776%22%5D%7D&r=1&s=1
- *S.1705 – AICT Act of 2021*: https://www.congress.gov/bill/117th-congress/senate-bill/1705/text?r=82&s=1

3

Deep Learning Deep Dive

In the age of AI implementation, the current period of AI we find ourselves in, we must understand the pros and cons of both **machine learning** (**ML**) and **deep learning** (**DL**) in order to decide when to use either technology. Some other terms you might have come across with respect to AI/ML tools are **applied AI** and **deep tech**. Again, both ML and DL are subsets of AI but operate at different levels of complexity and applicability. Traditional ML models are best suited for structured data (data that is labeled) and are able to make predictions or decisions without being explicitly programmed to do so. DL, on the other hand, represents an advanced evolution of ML that uses **artificial neural networks** (**ANNs**) to analyze and interpret much more complex, unstructured data.

We briefly touched on these concepts in *Chapter 1* but this distinction is important to mention here as well. With structured datasets, data scientists and ML engineers use the structure of the data to identify the key features that will be most relevant to the model. This process is known as **manual feature extraction**. With DL models, however, this process is done automatically by the ANN and is known as **automatic feature extraction**. This is because these models are capable of learning on their own, and this includes learning about which facets or "features" of an unlabeled (unstructured) dataset will make the biggest impact on the model.

As we've mentioned a few times in this book, the underlying tech that will, for the most part, power AI products will be ML or DL. That's because expert- or rule-based systems are slowly being powered by ML or not evolving at all. This is due in large part to the substantial value that incorporating ML can bring to the function of a legacy system. Also, all neural networks are referred to as "black box models" because of their opacity. As we covered in *Chapter 1*, one of the major distinctions between DL models and traditional ML models is their explainability. We still experiment with them tactically and try to understand them empirically and not so much theoretically. They still remain a black box and we try to understand them more and more through the actual use and practice of them. This will be largely true of all the models we will be discussing in this chapter.

So now, let's dive a bit further into DL and understand how the various models differ. More specifically, we will cover the following topics:

- Types of neural networks
- Exploring generative AI models

- Emerging technologies – ancillary and related tech
- Explainability – optimizing for ethics, caveats, and responsibility
- Guidelines for success

Types of neural networks

We'd like to first turn your attention toward some of the most popular kinds of neural networks used in DL today. Some of these may sound familiar, but it will help to familiarize yourself with some of these concepts, especially if you plan on working as a **product manager** (**PM**) for a DL product. Even if you aren't currently working in this capacity, you'll want to take a look through these in case your career does veer toward DL products in the future.

When considering using neural networks, an AI PM must weigh a variety of DevOps implications. DL models require high compute power, storage, data retrieval, and specialized skills and expertise. Hardware and infrastructure costs can be high and it can be challenging to see a strong ROI, particularly early on. In many cases, many will be faced with the decision to buy a neural network solution or build one in-house, which can also come with significant time investments as well that could delay time to market. These are all strategic decisions the AI PM will ensure get made and are all important considerations when building a product that aligns with the company's mission and resources.

In the following sections, we will touch on some of the most used ANNs in DL to give you an idea of what they are best suited for. As we did in the previous chapter with ML algorithms, we will describe some of the most popular use cases of each type of ANN so that you can understand, at least in a general sense, what some of the core competencies of each ANN are so that you can keep those ideas in mind should you pursue the creation of your own DL products in the future. If your aim is to specialize exclusively in supporting or building DL products of your own, this will be a summary overview of each ANN.

Multilayer perceptrons

After David Rumelhart, Geoffrey Hinton, and Ronald Williams's paper titled *Learning representations by back-propagating errors* came out in 1986, **multilayer perceptrons** (**MLPs**) were popularized because, in that paper, they used backpropagation to train an MLP. Unlike RNNs, MLPs are another form of feedforward neural network that uses backpropagation to optimize the weights. For this reason, you can think of MLPs as some of the most basic forms of ANNs because they were among the first to appear, and today, they're still used often to deal with the high compute power that's needed by some of the newer ANNs out there. Their accessibility and reliability are still useful today, which is why we wanted to start this list with MLPs to give us a good foundation for conceptualizing the rest of the DL algorithms.

The way they learn is the algorithm will send data forward through the input and middle layers to the output layer. Then, based on the results in the output layer, it will calculate the error to assess how off it was at predicting values. This is where backpropagation comes in because it will get a sense of how wrong it was in order to then backpropagate the rate of error. It will then optimize itself to minimize that error by adjusting the weights in the network and will effectively update itself.

The idea is you would pass these steps through the model multiple times until you were satisfied with the performance. Remember the distinction between supervised and unsupervised learning in *Chapter 1*? Because MLPs use backpropagation to minimize their error rate by adjusting weights, MLPs are a supervised DL algorithm as they know, based on the label data, exactly how far off they were from being right. These algorithms are also heavily used in ensembles with other ANNs as a final polishing stage because of their ability to handle complex, non-linear relationships in tabular data. Business use cases are most appropriate in classification problems like customer segmentation and personalization, fraud detection, forecasting, prediction, and scoring.

Case study

A financial services firm wants to update its credit scoring system because the existing model is based on traditional statistical models and is quickly becoming outdated. Given the influx of data and complexity in the market, their old system has become less effective. The firm is looking for a more sophisticated model that would handle the large volumes of data and the various data types and sources, and could also improve on predicting credit risk. If the new model could reduce the rate of defaults (primary goal) and improve creditworthy customer approval rates (secondary goal), the investment would be successful. Here's why an MLP was chosen as the favorite among the DL models tested by the AI PM team:

- MLPs can integrate easily with existing systems and handle complex, diverse datasets. So, it was a practical choice because it was integrated smoothly with the firm's existing infrastructure.
- MLPs excel at modeling non-linear interactions between features and this made it useful for capturing the complexities of a credit risk.
- The firm had diverse data that included transaction histories, credit behavior, demographic data, and appended external data sources, which meant they had to invest in a DL model that could process high-dimensional data elegantly. The MLP proved versatile with the data and effective at identifying high-risk customers, which addressed the primary goal.

The following is a list of other DL models that were tested but ultimately not chosen:

- CNNs: Because the firm's data was primarily tabular data that included transaction records and credit scores, a CNN would not have been a great choice.
- RNNs: These models are better suited for time series or sequential data so it wasn't fully able to capture the complexity of the data the firm had.
- Autoencoders: These models helped the firm identify relevant features in their high-dimensional data but, ultimately, the model helped point out associations more than it was able to classify someone as a credit risk. Experimentation with autoencoders was, however, useful once they landed on MLP.

After the deployment of its new DL-powered system, the firm saw an 18% reduction in default rates and a 10% increase in approval rates for creditworthy customers. This meant that both the primary and secondary goals of the product team were met.

Radial basis function networks

Radial basis function networks (**RBFNs**) came on the scene in 1988 with D.S. Broomhead and David Lowe's paper titled *Multivariable Functional Interpolation and Adaptive Networks*. RBFNs differ from most other ANNs we will cover in this chapter in that they only have three layers. While most ANNs, including the MLPs we discussed in the preceding section, will have an input and output layer with several hidden layers in between, RBFNs only have *one* hidden layer. Another key difference is that rather than having the input layer be a computational layer, the input layer only passes data to the hidden layer in RBFNs, so this ANN is incredibly fast. These DL algorithms are feedforward models, so they are computationally only really passing through two layers: the hidden layer and the output layer.

It would be helpful to think of these networks as similar to the KNN algorithm we discussed in the previous chapter, which aims to predict data points based on the data points around the value they're trying to predict. The reason for this is that RBFNs look to approximate values based on the distance, radius, or Euclidean distance between points and they will cluster or group data in circles or spheres to better make sense of a complex multivariable dataset similar to how a K-means clustering algorithm from *Chapter 1* would. This is a highly versatile algorithm that can be used with both classification and regression problems in both supervised and unsupervised ways.

Self-organizing maps

Self-organizing maps (**SOMs**) were introduced in the 1980s by Tuevo Hohonen and are another example of unsupervised competitive learning ANNs in which the algorithm takes a multivariable dataset and reduces it into a two-dimensional "map." Each node will compete with the others to decide whether it's the one that should be activated, so it's essentially just a massive competition, which is how it self-organizes. Structurally, though, SOMs are very different from most ANNs. There's just one layer or node outside of the input layer, which is called the Kohonen layer. The nodes themselves are also not connected the way they are in more traditional ANNs.

The training of a SOM mirrors our own brain's ability to self-organize and map inputs. When we sense certain inputs, our brain organizes those inputs into certain areas that are apt for what we're seeing, hearing, feeling, smelling, or tasting. The SOM will similarly cluster data points into certain groupings. The way that happens is through a learning/training process where the algorithm sends out the data through the input layer and weights, randomly selecting input data points to test against the nodes until a node is chosen based on the distance between it and the data point, which then updates the weight of the node. This process is repeated until the training set is complete and the optimal nodes have been selected.

SOMs will also be in the same class of clustering algorithms such as K-means, or the RBFNs we touched on in the preceding section, in that they are useful for finding relationships and groupings in datasets that are unlabeled or undiscovered.

Convolutional neural networks

Convolutional neural networks (**CNNs**), sometimes referred to as **ConvNets**, have multiple layers that are used largely for supervised learning use cases in which they detect objects, process images, and detect anomalies in medical and satellite images. The way this ANN works is through feedforward propagation, so it starts from the input layer and makes its way through the hidden layers to the ultimate output layer to categorize images. This type of ANN is characterized as categorical, so its ultimate goal is to put images into buckets of categories. Then, once they are categorized, it looks to group images by the similarities they share so that it can ultimately perform the object recognition that's used to detect faces, animals, plants, or signs on the street.

The four important layers in CNNs are as follows:

- The convolution layer: The convolution layer turns an image into a matrix of pixel values that are 0s and 1s and then further reduces that matrix into a smaller matrix that's a derivative from the first.
- The **rectified linear unit** (**ReLU**) layer: The ReLU layer effectively pares down the dimensions of the image that you pass to the CNN. Even color images are passed through a grayscale when they're originally assigned 0s and 1s. So, in the ReLU stage, the CNN actually gets rid of black pixels from the image so that it can reduce the image further and make it computationally easier for the model to process it.
- The pooling layer: This is another layer that reduces the dimensions of the image in another way. While the ReLU layer pares down the gradient in the image itself, the pooling layer pares down the features of the image, so if we pass the CNN an image of a cat, the pooling layer is where we will see various features such as the ears, eyes, nose, and whiskers identified. You can think of the convolution, ReLU, and pooling layers as operations that take segments of each image you feed your model and concurrently fire the outputs of those prior steps as inputs into the fully connected layer, which is what actually passes through the neural network itself to classify the image. In essence, the convolution, ReLU, and pooling layers prepare the image to pass through the neural network to arrive at a conclusion.
- The fully connected layer: This layer serves as a final step where high-level, learned features are converted into output predictions. This layer has a set of weights and biases that affect neurons from every prior layer, so it serves as the decision-making layer of the CNN.

CNNs can be used for things such as facial recognition, object identification, and self-driving cars or what's commonly referred to as computer vision applications of AI.

Recurrent neural networks

There are several operations that feedforward neural networks weren't able to do very well, including working with sequential data that is rooted in time, operations that need to contextualize multiple inputs (not just the current input), and operations that require memorization from previous inputs. For these reasons, the main draw of **recurrent neural networks** (**RNNs**) is the internal memory they possess that allows them to perform and remember the kind of robust operations required of conversational AIs such as Apple's Siri.

In contrast to the preceding CNN example, which works with a feedforward function, RNNs work in loops. Rather than the motion going from the input layer through the hidden layers and, ultimately, to the output layer, the RNN cycles through a loop back and forth, and this is how it retains its short-term memory. This means the data passes through the input layer, then loops through the hidden layers, before it ultimately passes to the output layer. It's important to note that RNNs only have short-term memory, which is why there was a need for an LSTM network. More on that in the next section.

In essence, the RNN actually has two inputs:

- The first is the initial data that makes its way through the neural network.
- The second is actually the information and context it's acquired along the way.

This workflow is the framework with which it also effectively alters its own weights to current and previous inputs, so it's course-correcting as it goes through its loops. This process of retroactively adjusting weights to minimize its error rate is known as backpropagation, which you'll recall from *Chapter 1* (*A brief history of DL*) as this was a major advancement that has helped DL become so successful.

It's helpful to imagine that an RNN is actually a collection of neural networks that are continuously retrained and optimized for accuracy through backpropagation, which is why it's also considered a supervised learning algorithm. Because it's such a robust and powerful DL algorithm, we can see RNNs used for anything from captioning and understanding images to predicting time-series problems to natural language processing and machine translation.

RNNs do well with sequential data and place a premium on the context in order to excel at working with time-series data, DNA and genomics data, speech recognition, and speech-to-text functions.

Long short-term memory networks

Long short-term memory (**LSTM**) networks are basically RNNs with more memory power. Often, the way they manifest is through networks of the LSTM because what they do is connect layers of RNNs, which is what allows them to retain inputs over lags or longer periods of time. Much like a computer, LSTMs can write, read, or delete data from the memory they possess. Because of this, they have the ability to learn about what data they need to hold onto over time. Just as RNNs continuously adjust their weights and optimize for performance, LSTMs do the same thing by assigning levels of importance for what data to store or delete through their own weights.

LSTMs mirror our own ability to discard irrelevant or trivial information through time through LSTM cells, which have the ability to let the information come in as an input, be forgotten or excluded completely, or let it pass to influence the output. These categorizations are referred to as gates and they're what allow LSTMs to learn through backpropagation.

Deep belief networks

Deep belief networks (**DBNs**) also have multiple layers, including multiple hidden layers, but the nodes in one layer aren't connected to each other, though they are connected to nodes in other layers. There are relationships between the layers themselves, but not between the nodes within. DBNs are unsupervised learning layers of what are called **restricted Boltzmann machines** (**RBMs**), which are themselves another form of ANN. These layers of RBMs are chained together to form a DBN. Because of this chain, as data passes through the input layer of each RBM, the DBN learns and obtains features from the prior layer. The more layers of RBMs you add, the greater the improvement and training of the DBN overall. Also, every RBM is taught individually and the DBN training isn't done until all the DBNs have been trained.

DBNs are referred to as **generative ANNs** because each of the RBMs learns and obtains potential values for your data points based on probability. Because of this generative ability that they have, they can be used for things such as image recognition, capturing motion data, or recognizing speech. They are also computationally energy-efficient because each cluster of RBMs operates independently. Rather than data passing through all the layers in concert as with feedforward ANNs, data stays local to each cluster.

> As a product manager, you won't need to have in-depth knowledge of each neural network because, if you're building a product with DL components, you've got internal experts who can help determine which neural networks to use. But it does help to know what some of the most common types of neural networks out there are so that you aren't left in the dark about those determinations.

In the next section, we will see how DL neural networks relate to generative AI.

Exploring generative AI models

In the previous section, we went over the majority of the ANNs that make up DL models. In this section, we will go deeper into the risks and rewards of various generative AI models specifically, given the exposure and potential generative AI has in the world of AI products. The world of generative AI is still being formed and the below models offer a snapshot of the main model types that make up the current generative AI landscape.

Keep in mind that the beauty of generative models is that their outputs will differ every time you prompt them. Some will be better at predicting what you're expecting than others and some will be able to give you more diversity than others. What you will get by the end, however, will be that model's best guess for meeting your request. Because of a variety of probabilistic and deterministic decisions the model is making across each layer of the neural network, the outputs we get from generative models are always novel and never exactly the same.

Because of their flexibility across applications and use cases, generative AI models are used for everything from facial recognition, data augmentation and synthesis, anomaly detection, super-resolution imaging and enhancement, and content creation, whether that be video, sound, image, or text. In the following sections, we will explore some of the major categories of generative AI models that are widely used today, as well as some of their strengths and drawbacks.

Generative adversarial networks

Generative adversarial networks (**GANs**) are my favorite type of ANN because they're essentially made up of two neural networks (the generator and the discriminator) that are pitted against each other, hence the name, which compete toward the goal of generating new data that's passable for real-world data. Because of this generative ability, GANs are used for image, video, and voice generation. They were also initially used for unsupervised learning because of their generative and self-regulation abilities, but they can be used for supervised and reinforcement learning as well. The way they work is one of the neural networks is referred to as the generator and the other is the discriminator, and the two compete as part of this generative process.

GANs were first introduced in a breakthrough paper that came out in 2014 by Goodfellow et al. titled *Generative adversarial networks*, which states that GANs "simultaneously train two models: a generative model G that captures the data distribution, and a discriminative model D that estimates the probability that a sample came from the training data rather than G. The training procedure for G is to maximize the probability of D making a mistake."

We can think of discriminative and generative models as two sides of the same coin. Discriminative models look at the features a type of image might have, for example, looking for associations between all the images of dogs the model is currently learning from. Generative models start from the category itself and expand out into the potential features that a category in that image might possess. If we take the example of a category such as space kittens, the generative model might look at the example data it's fed and deduce that if it creates an image, it should create something that involves space and kittens. The discriminative model then takes the image the generative model creates and confirms, based on its own learning, that any images in the space kittens category must contain both kittens and space as features.

Another way to explain this is that the generative model maps the label to potential features and the discriminative model maps features to the label. What's most interesting to us about GANs is they effectively pass or fail their own version of the Turing test. How do you know whether you passed? If the GAN correctly identifies a generated image as a falsified image, it's passed (or failed?) its own test. It really depends on how you look at passing or failing for that matter. If it incorrectly labeled a falsified/generated image as a "real" image, it means the generative model is pretty strong because its own discriminator wasn't able to discriminate properly. Then again, because it's a double-sided coin, it means that the discriminator needs to be strengthened to be more discerning. GANs are very meta.

The steps a GAN takes to run through its process are as follows:

1. It begins with a generator neural network that takes in data and returns an image.

2. The image is then fed to the discriminator along with other images from a real-world dataset.
3. Then, the discriminator produces outputs that are numbered between 0 and 1, which it assigns as probabilities for each of the images it is discriminating, with 0 representing a fake and 1 representing an authentic real-world image.
4. GANs also use backpropagation, so every time the discriminator makes a wrong call, the GAN learns from previous mistakes to correct its weights and optimize itself for accuracy.

Because of the adversarial nature of GANs, the effects of pitting two neural networks against each other offer us a tool that produces the most realistic generated images. But it's not just images GANs are good at making; they're also great for producing synthetic data. Neural networks need so much data that, often, creating synthetic data is the best way to adequately feed them, but the drawback is if you create enough synthetic data, the model will collapse upon itself. GANs also excel at translating between images, so you often see them used in map generation from satellite images or in products that allow you to edit artistic styles, filter between images, or even improve the resolution of blurry images. Because they're so good at producing outputs that are diverse, they've also been used to improve **natural language processing** (**NLP**). The superpower of GANs is the diversity of images, text, and video that can be derived from using them.

> In the context of DL, synthetic data refers to data that is artificially generated rather than collected from real-world events and observations. It's valuable for training, validating, and testing DL models, where obtaining enough real-world data is challenging, time-consuming, or costly.

However, GANs can suffer, like any ANN, if they don't have diverse enough data samples to learn from. Generated images are often flawed and you can see these flaws anytime a generated image adds a sixth finger to someone's hand. These models are trying to predict the output the user is asking for based on the prompt it's given, so we often see the limitations of generated images or video when the output is too flawed to be real. They're also hard to train and can be unstable. Ethically, because GANs are often so focused on image generation, and because that image generation is so crisp, they can be used for malicious intent by those looking to influence the political, economic, social, or cultural landscape through the use of propaganda and misinformation.

Autoencoders

Autoencoder models are the simplest of the generative AI models we will cover as they primarily exist to reduce noise in data and to learn how to reconstruct basic patterns in the data they see. For that reason, it's used most for unsupervised learning tasks like data compression, denoising, and feature extraction to identify the most important features in a data sample, which means it's often used with other models as part of a larger model pipeline. Let's break down the layers:

- The first part of autoencoders is called an **encoder**. This part of the model has an input layer where inputs flow into. Those inputs flow through two hidden layers that reduce the dimensionality of the data. The encoder stage does exactly what its name suggests; it turns the input data it receives into code that will pass through to the second part.

- The second part is the decoder, which "decodes" and reconstructs an output that's as close to the original input as it was received but with a few differences. This part also has two hidden layers in between the inputs it receives from the encoder and what it ultimately produces as an output.

Autoencoders are especially good at finding patterns in your data, removing noise from that data, and producing results based on the patterns they see. Note that, unlike GANs, where you use two neural networks against each other, these models operate as one unit.

Over time, after passing inputs through the layers again and again, they learn to differentiate what's important and relevant about an image and what isn't. In terms of functionality, autoencoders themselves do not compete with GANs, but they may be used to produce more data or to reduce noise from the training data for the generator network within GANs, making them more effective. The main thing to note here is the *latent space*, where input data is mapped and encoded, is deterministic in autoencoders. This means the inputs are fixed in the space. If the latent space were instead probabilistic, that would then open up a whole new world of possibilities. This would also turn an autoencoder model into a variational autoencoder.

Variational autoencoders (**VAEs**) take this same principle of coding and decoding an image or text and they add a twist to this process by turning that same input data an autoencoder would have and mapping it into a range of probabilities rather than just one fixed spot. What this does is it opens the model to quite a large variety of potential outputs to choose from in comparison with autoencoder models. Because of this probabilistic quality of VAEs, they become more useful for generative purposes and compete with GANs in terms of offering generated outputs that are diverse and novel enough. Training VAEs is a more demanding process than it is for autoencoders. For autoencoders, the task is simple: they exist to reduce noise and work with fixed points in the latent space. The process of learning and applying those lessons across a wide range of probabilities requires more compute power but they're still easier than training GANs, which is why they're such a close competitor to them.

Some of the drawbacks of autoencoders relate to images being blurry or not being realistic enough in comparison to GANs, sacrificing specificity for diversity. These models are also susceptible to collapse if they don't see enough diverse examples in their training data, resulting in outputs that start to resemble each other too closely. As with all model types, you're optimizing for certain conditions based on the training data they're learning from. Depending on what that data is, one model might work better than another and the fastest path to learning which works best is to simply pass the data through different model types.

Diffusion models

Many of the most popular generative AI applications out there (from Open AI's DALL-E 3 to Midjourney to Meta's Emu, Google's Imagen, Stability.ai's SDXL, and Adobe's Firefly) are all based on diffusion models. The reason they're called diffusion models is they're effectively dissolving training data into noise, learning from that process, and using that learning to reverse engineer its wisdom to create totally new images. If you can think of static on a TV as "noise," that static represents random potentiality. But an image or video is not random; it is specific. So in diffusion models, the slow process of progressively adding more and more static to an image until it's indetectable is the way it understands the qualities that give the image its training on its specificity.

> In a paper written by Open AI's Prafulla Dhariwal and Alex Nichol titled *Diffusion models beat GANs on image synthesis*, the authors attribute a number of advantages to diffusion models and clue us into a big part of why they've been so successful in saying "GANs currently hold the state-of-the-art on most image generation tasks as measured by sample quality metrics such as FID, Inception Score and Precision. However, some of these metrics do not fully capture diversity, and it has been shown that GANs capture less diversity than state-of-the-art likelihood-based models (diffusion models). Furthermore, GANs are often difficult to train, collapsing without carefully selected hyperparameters and regularizers. While GANs hold the state-of-the-art, their drawbacks make them difficult to scale and apply to new domains. As a result, much work has been done to achieve GAN-like sample quality with likelihood-based models. While these models capture more diversity and are typically easier to scale and train than GANs, they still fall short in terms of visual fidelity. Furthermore, except for VAEs, sampling from these models is slower than GANs in terms of wall-clock time."

We see a number of advantages in the passage above ranging from training difficulty, output diversity, model collapse, and scalability. They do, however, go on to say that although they "believe diffusion models are an extremely promising direction for generative modeling, they are still slower than GANs at sampling time due to the use of multiple denoising steps." Depending on the size of the diffusion model, speed will be an issue as each denoising step can take anywhere from five to twenty times longer to train than GANs.

Transformer models

While many of the **large language models** (**LLMs**) have been created using other ANNs like RNNs and LSTMs, which we referenced in the previous section of this chapter, recent advancements in this space have come from transformer models. Transformers have become the darlings of the NLP space due, in large part, to their ability to handle sequential data so well. As you can imagine, a model that can holistically contextualize, understand, and respond to a prompt is quite the breakthrough when it comes to language understanding. Popular transformer models include Stable Diffusion, GitHub's Copilot, Quora's Poe, Microsoft's MS Copilot, Anthropic's Claude, Google's BERT, T5, and Gemini models, as well as OpenAI's GPT model.

First introduced in 2017 by a team of Google researchers in a paper titled *Attention is all you need*, transformers made it possible for a model to analyze the entirety of a body of text rather than trying to understand it sequentially. What that means is that rather than understanding each word in order, a transformer model allows for a more global understanding of an entire prompt. It's able to do this by assigning weights or various levels of "attention" to different words in the text it's trying to analyze or deliver. It does this through a process called parallelization.

Another major advancement of the transformer model is the ability to train on huge datasets in comparison to RNNs or LSTM, while also being able to support such a large number of parameters. The world's current state-of-the-art LLM, GPT-4, is owned by OpenAI and it currently has 1.8 trillion parameters. You can think of parameters as individual "lessons" or variables the model learns from its training data. Note this is not the same as hyperparameters, which are a different set of tools that ML engineers can use to calibrate or fine-tune their models.

While these models are incredibly powerful, they are quite expensive to build, train, run, and keep updated as they require significant compute power for all these aspects. Because transformer models understand and answer at such a global level, they also require a lot of memory to function. The main drawback of these models is the ethical repercussions of using outputs from language models as fact. While they often do deliver factual information, they may make some up in the form of "hallucinations," a term actively used to describe falsities generated by LLMs.

While it may seem logical to question everything that's created by generative models, people may not always respond to generated content that way. They may, in many cases, take for granted the authenticity or appropriateness of generated content they come across, whether it's text, video, or images. Even if people do know that what they're seeing may not be authentic or appropriate, over time and with added exposure to certain kinds of content, they may still be influenced in ways that change their behavior or attitude toward certain things.

Emerging technologies – ancillary and related tech

ML and DL have been used heavily in applications related to NLP, speech recognition, chatbots, virtual agents and assistants, decision management, process automation, text analytics, biometrics, cybersecurity, content creation, image and emotion recognition, and marketing automation. It's important to remember, particularly from a PM's perspective, that AI will increasingly work its way into more of how we live our lives and do our work. This is doubly true if you work in an innovative capacity as a PM where you're involved with the ideation and creation of new use cases and **mimimum viable products** (**MVPs**) for future products.

Over the passage of time, we'll see AI continue to augment our workforce both through the process of internal automation as well as through the adoption of AI-based no-code/low-code external products and applications that will boost job functions, skills, and abilities across the board. AR, VR, and the metaverse also offer us new emerging fields where ML will learn more about our world, help us learn about ourselves, and also help us build new worlds altogether. We will also continue to see ML employed through AI-powered devices such as self-driving planes, trains, and automobiles, as well as biometrics, nanotechnologies, and IoT devices that share streams of data about our bodies and appliances that we can increasingly use to optimize our security, health, and energy usage.

Blockchain technologies offer a lot of opportunities to address authenticity and provenance, particularly in the context of text, image, and video generation. Since generative AI models can create highly realistic content, we have to have a way of confirming what was artificially generated. Ensuring that outputs are tracked back to the source model that created them is a way we can maintain credibility and prevent the misuse of AI-generated content. A blockchain can log the generation process of media outputs and offer a record of how, when, and by which model content was created. As time goes on and generative AI models become more enmeshed in our reality, it will be important for us to be able to confidently say which images are potential forgeries or unauthorized reproductions.

Blockchain's decentralized nature also supports authenticity because metadata regarding generated content would be stored on the blockchain, so being able to prove that content was generated by a trusted, licensed model is a way to help protect content creators as well as consumers of content. This means that transparency and accountability can be effectively mandated by the blockchain because it could require platforms that use generative models to have a unique blockchain ID to link to its generation history. This creates more transparency around ownership and helps keep licensing fair as well. Creators can store proof of ownership, licenses, and usage permissions on the blockchain to safeguard against theft and misuse.

Quantum computing can massively help us deal with the ongoing issue of storing and retrieving data, particularly big data, in cost-effective ways. As recently as December 2023, IBM released the first quantum computer with more than 1,000 qubits, called Condor. This has significant implications for DL and AI development because quantum computers like Condor can help DL models learn faster or even allow for more complex ANNs that might have been previously inaccessible due to computational constraints. Because so many DL projects can take weeks to train and require access to big data, ancillary developments in quantum computing can prove groundbreaking in the area of DL to the point where the algorithms both require less data to train on and can also handle more data and compute power more quickly.

As the years go on, we will likely see AI play an increasingly vital role in enabling efficiency, security, and user experience for Web3 technologies like **decentralized autonomous organizations** (**DAOs**), governance, **decentralized finance** (**DeFi**), smart contracts, personalized **non-fungible tokens** (**NFTs**), and metaverse development. DL algorithms can analyze large datasets generated by DAO activities and enhance security protocols by helping to identify trends. They can also help generate NFTs or generate more intelligent **non-player characters** (**NPCs**) for the metaverse. Other generative models focused on image, sound, and video generation can further enhance our experience of Web3 capabilities like AR/VR virtual world creations in the multiverse as well. The main point to highlight here is that generative AI models are just especially good at mimicking outputs that are digestible by the human senses: things we can certainly see, hear, touch, and potentially even taste and smell.

Explainability – optimizing for ethics, caveats, and responsibility

Ethics and responsibility play a foundational role in dealing with your customers' data and behavior, and because most of you will build products that help humans make decisions, eventually, someone is going to ask you how your product arrives at conclusions. Critical thinking is one of the foundational cornerstones of human reasoning and if your product is rooted in DL, your answer won't be able to truly satisfy anyone's skepticism. Our heartfelt advice is this: don't create a product that will harm people, get you sued, or pose a risk to your business.

That said, the goal post of what constitutes “harm” or an “AI system” is always changing and will continue to change, as we saw with the current regulations and movements to legislate, which were introduced in *Chapter 2* of this book. Companies, technologists, leaders, and officers of organizations often won’t be held to scrutiny or accountability until after harm has been done. Not everyone operates in the world with a moral compass, and many may support the building of products they’re aware cause harm. Many will remain willfully ignorant as a matter of choice.

If you’re leveraging ML or DL in a capacity that has even the potential to cause harm to others, if there’s a clear bias that affects underrepresented or minority groups (in terms of race, gender, or culture), go back to the ideation phase. This is true whether that’s immediate or downstream harm. This is a general risk that ML poses to us collectively: the notion that we’re coding our societal biases into AI without taking the necessary precautions to make sure the data we feed our algorithms is truly unbiased.

The engineers who build these ANNs are unable to look under the hood and truly explain how ANNs make decisions. As we’ve seen with the, albeit layman, preceding explanations of DL algorithms, ANN structures are built on existing ML algorithms and scaled, so it’s almost impossible for anyone to truly explain how these networks come to certain conclusions. Nevertheless, some of the most encouraging legislation brought forth by the EU, the AI Act discussed in *Chapter 2*, aims to combat this very issue. By insisting that companies not only demonstrate the level of risk their AI systems pose but also prove their findings in the form of a white paper that aims to specifically articulate the potential harm, they’re combatting the main issue that plagues DL models. Going a step further, they also want there to be requirements that communicate the training methodology used for the creation of an AI system, which can allow for greater trust and transparency in AI as it evolves. This is encouraging to see and a positive move forward when it comes to the promotion of AI regulation and safety.

DL algorithms are often referred to as a black box because they resist a truly in-depth analysis of the underpinnings of the logic that makes them work. This is true of all deep tech models, including the most popular and effective generative AI models. DL has a natural opacity to it because of the nature and complexity of ANNs. Remember that ANNs effectively just make slight adjustments to the weights that affect each neuron within its layers. They are basically elaborate pattern finders using math and statistics to make optimizations to their weighting system. They do that hundreds of times for each data point over multiple iterations of training. We simply don’t have the mental capacity or language to explain this.

You also don’t have to be a DL engineer to truly understand how your product affects others. If you, as a PM, are not able to fully articulate how DL is leveraged in your product and, at the very least, can’t prove that the outputs of your DL product aren’t causing harm to others, then it probably isn’t a product you want to go all in for. Now that we’ve seen DL become more accessible through the use of popularized generative AI models, we see the potential for the opacity of DL to become exacerbated and fuel even more distrust in the market. While DL is currently more accessible, it’s still very much in the research phase and many PMs should be prudent when incorporating it because of this issue of explainability. We urge PMs to use caution when looking to wet their feet with DL. Moving forward without a sense of stewardship of and responsibility for our peers and customers with something as complicated and full of potential as DL is a recipe for adding more chaos and harm to the world.

AI PMs play a crucial role in making sure the AI systems they build are designed, deployed, and monitored with ethical considerations in mind. The following is a list of best practices that are meant to serve as a checklist for AI PMs to use to ensure that they've got their bases covered:

- Schedule periodic audits of your AI system to ensure certain groups are not discriminated against by addressing bias in training datasets and model predictions. Tools like IBM's Fairness 360 toolkit and guides like TensorFlow's Fairness Indicators can help with this.
- Set up tools to track predictions and user interactions so that there's a triggering system when there are unusual patterns that could be proxies for ethical concerns.
- Create accessible documentation for AI systems where users/developers can ask questions or raise concerns.
- Proactively ask for feedback through user surveys and feedback forms to get information on system performance and ethical concerns.
- Regularly review training data and model outputs for fairness and inclusivity.
- Document ethical considerations and decisions throughout the development lifecycle, particularly for key decisions, and use **explainable AI** (**XAI**) methods like LIME or SHAP to document model interpretability and decision-making. DataCamp has a great resource for that, which can be found here: `https://www.datacamp.com/tutorial/explainable-ai-understanding-and-trusting-machine-learning-models`.
- Regularly review changes in AI regulations and legislation with legal teams and establish an AI ethics review board within your organization to demonstrate due diligence regarding ethical concerns.
- If a cross-functional governance and ethics board doesn't exist at your organization, help form one to ensure leadership, technical, and legal stakeholders are represented where you can discuss strategic decisions together.
- Embed ethical principles into the design process from the outset to make sure **fairness, accountability, and transparency** (**FAT**) remain constants within your product team or organization. Adopt ethical frameworks like IEEE or OECD, which outline responsible AI development.

Do we always have to be so cautious? Not necessarily. If DL applications get really good at saving lives by detecting cancer or they work better when applied to robotics, who are we to stand in the way of progress? However, medical applications of DL models are heavily scrutinized for accuracy and safety standards. As time goes on, we will be seeing more legal, legislative, and regulatory oversight of AI systems, not less. The complexity and depth of DL models make it difficult to trace how certain decisions or predictions are made compared to traditional ML models. If you're an AI PM, you will have to help prepare your organization to account for this ambiguity.

Guidelines for success

When it comes to DL, we can only truly grapple with its performance. Even from a performance perspective, a lot of DL projects fail to give the results their own engineers are hoping for when everything goes "right," so it's important to manage expectations across functions. This includes managing the expectations of your leadership team as well. If you're an AI PM or entrepreneur and you're thinking of incorporating DL, do so in the spirit of science and curiosity. Remain open about your expectations.

Make sure you're setting your team up for success by focusing on data quality as well. A big part of your ANN's performance lies in the data preparation you take before you start pre-training your models. Passing your data through an ANN is the last step in your pipeline. If you don't have good validation or if the quality of your data is poor, you're not going to see positive results. Then, once you feel confident that you have enough data and that it's clean enough to pass through an ANN, start experimenting. If you're looking for optimal performance, you'll want to try a few different models, or a selection of models together, and compare the performance.

The time it takes to fine-tune a DL model is also aggressively long. The use of popularized or open source generative AI models out there through Hugging Face, OpenAI, Google Colab, and others could make this process shorter, but working with these tools still constitutes a big investment. Even if many tools offer free or low-cost options, running large models and training them would still require a lot of compute power, which would be a cost investment for cloud services. DL models can become a financial burden if not carefully managed.

Data acquisition, storage, and transformation are additional costs. Using Google Colab might technically be free for experimentation, but once you train extensively or deploy your experiment, you're still going to account for those costs. If you're used to other forms of ML, it might shock you to experience training a model over the course of days or weeks. This is largely because the amount of data you need to train ANNs is vast; most of the time, you need at least 10,000 data points, and all this data is passed through multiple layers of nodes and processed by your ANN. Your ANN is also, most of the time, going to be an ensemble of several types of ANNs we mentioned previously. The chain then becomes quite long.

The nature of ANNs is inherently mysterious because of the complexity of the layers of artificial neurons. We cannot see deterministic qualities. Even when you do everything "right" and you get a good performance, you don't really know why. You just know that it works. The same goes for when something does go wrong or when you see poor performance. Once again, you don't really know why. Perhaps the fault lies with the ANN or with the method you're using or something has changed in the environment. The process of getting back to better performance is also iterative. And then it's back to the drawing board.

Remember that these are emerging tech algorithms. It may take us all some time to adjust to new technologies and truly understand the power they have. Or not! Part of the disillusionment that's happened with DL actually lies in the tempering of expectations. Some DL algorithms can make amazing things happen and can show immensely promising performance but others can so easily fall flat. It's not a magic bullet. It's just a powerful tool that needs to be used in the proper way by people who have the knowledge, wisdom, and experience to do so. Considering most of the ANNs we went over together are from the 80s, 90s, and early 2000s, that's not much time.

Even generative AI models, for all their hype and buzz, can fall short of expectations. They're also expensive to make your own. To put things in perspective: you might spend upward of $700,000 to pre-train a language model with 70 billion parameters and 100 billion tokens through several epochs and still have a language model that doesn't work very well for your intended use case. So, tread carefully here if you're building, managing, or ideating on DL products. When in doubt, there are other more explainable models to choose from, which we covered in *Chapter 1*. It's better to be safe than sorry. If you've got lots of time, patience, excitement, and curiosity, along with a safe, recreational idea for applying DL, then you're probably in a good position to explore that passion and create something the world could use in good faith.

The following are a few guidelines you can use to make sure you're setting up your product organization for success:

- Set realistic goals and communicate that DL is experimental and iterative in nature. Make it clear that success will come from trial and error.
- Ensure robust data validation, cleaning, and preparation steps before passing data into ANNs, and make sure this is treated as one of the highest priorities before model training.
- Test multiple models or a combination of models and compare their performance to find the best fit for your use case. Use experimentation as a way to discover the right blend of model complexity and performance for your product.
- Be prepared for long development cycles and plan for resource-intensive computing costs. Free tools can only take you so far before you face compute and storage expenses. Bring this up early and often with your financial stakeholders.
- Budget for compute and storage costs and monitor expenses regularly to avoid runaway costs during experimentation and deployment.
- Use post hoc analysis tools like LIME or SHAP to help with explainability should it come up downstream.
- Approach DL with curiosity but don't overpromise. Stay grounded in reality and avoid hyping up your product or expected performance, particularly when managing expectations with stakeholders early on.
- Reach for more explainable ML models for critical applications where transparency is paramount. DL should only be used when its advantages outweigh the challenges of explainability.

Summary

We got the chance to go deep into DL in this chapter and understand some of the major influences that impact this subsection of ML. We also got the chance to look at some of the specific ANNs that are most commonly used in products powered by DL in order to get more familiar with the actual models we might come across as we build with DL. We ended the chapter with a look at some of the other emerging technologies that collaborate with DL, as well as getting further into some of the concepts that impact DL most: explainability and guidelines for success.

DL ANNs are super powerful and can exhibit great performance, but if you need to explain them, you will run into more issues than you would if you stick to more traditional ML models. We've now spent the first three chapters of the book getting familiar with the more technical side of AI product management. Now that we've got that foundation covered, we can spend some time contextualizing all this tech.

In the next chapter, we will explore some of the major areas of AI products we see on the market, as well as examples of the ethics and success factors that contribute most to commercialization.

References

- *Some methods for classification and analysis of multivariate observations* by J. McQueen
- *K-Nearest neighbors algorithm: Classification and regression star*: `https://www.historyofdatascience.com/k-nearest-neighbors-algorithm-classification-and-regression-star/`
- *Random forests*: `https://www.stat.berkeley.edu/~breiman/randomforest2001.pdf`
- *Decision trees*: `https://www.cse.unr.edu/~bebis/CS479/PaperPresentations/DecisionTrees.pdf`
- *Support vector machine*: The most popular machine learning algorithm: `https://cml.rhul.ac.uk/svm.html`
- *Logistic regression*: `https://uc-r.github.io/logistic_regression`
- *Logistic regression history*: `https://holypython.com/log-reg/logistic-regression-history/`
- *Bayes Classifier*: `https://www.sciencedirect.com/topics/computer-science/bayes-classifier`
- *Principal component analysis*: `https://www.sciencedirect.com/topics/agricultural-and-biological-sciences/principal-component-analysis#:~:text=PCA%20was%20invented%20in%201901,the%20modeling%20of%20response%20data`
- *Galton, Pearson, and the peas*: A brief history of linear regression for statistics instructors: `https://www.tandfonline.com/doi/full/10.1080/10691898.2001.11910537`
- *IBM promises 1000-qubit quantum computer—a milestone—by 2023*: `https://www.science.org/content/article/ibm-promises-1000-qubit-quantum-computer-milestone-2023`
- *Boston Dynamics says AI advances for Spot the robo-dog are coming*: `https://venturebeat.com/ai/boston-dynamics-says-ai-advances-for-spot-the-robo-dog-are-coming/`
- *Convolutional neural network tutorial*: `https://www.simplilearn.com/tutorials/deep-learning-tutorial/convolutional-neural-network`
- *Generative adversarial networks*: `https://arxiv.org/abs/1406.2661`
- *The self-organizing map*: `https://sci2s.ugr.es/keel/pdf/algorithm/articulo/1990-Kohonen-PIEEE.pdf`
- *Multivariable functional interpolation and adaptive networks*: `https://sci2s.ugr.es/keel/pdf/algorithm/articulo/1988-Broomhead-CS.pdf`
- *Long short-term memory*: `http://www.bioinf.jku.at/publications/older/2604.pdf`

- *Learning representations by back propagating errors*: `https://www.iro.umontreal.ca/~vincentp/ift3395/lectures/backprop_old.pdf`
- *The numerical solution of variational problems*: `https://www.sciencedirect.com/science/article/pii/0022247X62900045`
- *The perceptron: A perceiving and recognizing automaton*: `https://websites.umass.edu/brain-wars/1957-the-birth-of-cognitive-science/the-perceptron-a-perceiving-and-recognizing-automaton/`
- *A logical calculus of the ideas immanent in nervous activity*: `https://www.cs.cmu.edu/~./epxing/Class/10715/reading/McCulloch.and.Pitts.pdf`
- *AI education*: `https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report-_Chapter-4.pdf`
- *Diffusion models beat GANs on image synthesis*: `https://arxiv.org/abs/2105.05233`
- *Attention is all you need*: `https://arxiv.org/pdf/1706.03762`

4

Commercializing AI Products

Now that we're in the period of **artificial intelligence** (**AI**) integration, we're seeing many use cases of AI proliferating across industries. In our work managing AI products, we've certainly relied on AI consultants and PhD-level advisors to help us with modeling and orchestrating our data strategy to support a full-scale AI operation. However, as the rising tidal wave of AI continues to penetrate various companies and use cases, we're seeing less of a reliance on breakthroughs achieved through advanced degrees. What's most important now is familiarity with the use of even simple, reliable models. There's a time and place for specialization. Data science and AI are massive umbrella terms but, based on my experiences as a product manager, I see a real need for data and AI generalists who can understand the use cases themselves and how they relate to a business perspective.

In the age of AI breakthroughs pre-smartphones, we didn't have a wide selection of use cases because the data wasn't as abundant as it is now. The rise of generative AI models has also helped with regard to the availability of data as many ANNs require so much training data to train on and optimize those models. In many cases, neural networks are used to create synthetic data as well. Now that data is abundant, we're seeing a profound shift – instead of the focus being on researching and discovering the most cutting-edge algorithms, it is on acquiring enough talent to use tried-and-true models and algorithms.

What we particularly like about this is the concept of AI democratization, or the idea that AI is becoming accessible to more groups of people outside research and big tech companies. This is especially true in 2024 as generative AI models are sweeping the landscape in terms of use cases and applications. The ChatGPT API alone is now responsible for the proliferation of so many generative applications that are out there in the market today.

As a product manager, your philosophy on this will really depend on who you ask. Some product managers believe that the way to create the best product is to conduct your own breakthrough research from an AI perspective. Other product managers believe you can actually get a lot done with models that have become commonplace or have been around a long time, such as regression models. This remains true and even more relevant today where generative capabilities are concerned. The reason for this is training the kind of complex ANNs that generate content can be so prohibitively expensive that for the time being, we will see the market open up to models that have been trained by heavily financed, open sourced organizations.

The goal of a product manager is to help the manifestation of a product that isn't just a commercial success but an industrial triumph as well. What exactly do we mean by industrial triumph? If we can influence you to bring a product to market that truly can offer something new, necessary, and useful to the world, we will have achieved the central goal of writing this book. We would also consider that a triumph.

Explainability is a big area of interest for a product manager. The idea that you can easily explain your AI product, how it comes to certain conclusions, and the inner mechanics of how it works becomes very important when you're dealing with the public or a sensitive subject matter, or if you have high-profile or enterprise customers that may be subject to risk later downstream. We don't want to convince you of what kind of product manager you should be; rather, we want to lay down the options for you to consider.

The purpose of this chapter is to highlight some promising developments in AI products in recent years to offer a variety of perspectives on how AI can be leveraged across a number of different business models. This isn't about emulating other companies but rather coming to appreciate how AI philosophy, behaviors, ethics, and practices can change from one business model to another. For us, this is vital to how the product is managed at the department level. As a product manager, you are tasked with the commercial success of the product, and understanding how AI relates to your business model is important to your risk management.

With that, let's explore the best examples of AI product management done well and see how this success impacts a business in all areas, expected and unexpected. In this chapter, we will be taking a look at some examples of AI companies through the lens of differing business models and offerings, including **business-to-business** (**B2B**) examples and **business-to-consumer** (**B2C**) examples. We will also be taking a look at companies thriving in unsaturated and underserved markets (blue oceans) as well as saturated and competitive markets (red oceans). The aim of this chapter is to highlight a few key areas in which businesses are thriving with regard to the AI products they bring to market.

The chapter will cover the following topics:

- The professionals – examples of B2B products done right
- The artists – examples of B2C products done right
- The pioneers – examples of blue ocean products
- The rebels – examples of red ocean products
- The GOATs – examples of differentiated disruptive and dominant strategy products

The professionals – examples of B2B products done right

The professionals are a good place to start with grouping products. AI will gravitate first toward use cases that can be profitable and allow for research and optimization. Because B2B products are products that are made for and used by other businesses, their use case is oriented completely toward the business world. This impacts everything from how they're marketed to how they're bought, sold, used, and negotiated. So many B2B products speak to the business impact that a product can satisfy for a customer across multiple levels. It's a great way to learn potentially helpful applications of AI.

Part of the challenge with the rising supply of AI companies is that they need data to train on. Specifically, one of the ethical challenges with this expansion in data and AI products is being able to offer large amounts of data without leaking information that can identify you as a person, also known as **private personal information** (**PPI**) or **personally identifiable information** (**PII**). Hazy, a UK-based AI company, offers its customers just that: the ability to derive insights, understand signals, and share data using synthetic data. Because of the nature of data-hungry deep learning models, synthetic data is preferred when training neural networks, and because of that, Hazy has a bright future ahead of itself. After an initial seed funding round of $3 million, it acquired $9 million in funding during a series A round in March 2023 led by Conviction VC. Their most recent round included funding from UCL Technology Fund and M12, a Microsoft venture capital fund.

The nature of a successful B2B company lies in its ability to create success for the businesses it supports, and though we're firmly in the data-rich era of big data, machine learning and deep learning still require massive volumes of data to learn and retrain from regularly. Hazy has done a fantastic job of alleviating businesses' pain points when it comes to data availability, and it's evident in the loyalty of their customer base. Another way Hazy maintains this loyalty is by educating its customers on the ethical and legal ramifications of traditional ways of anonymizing or masking existing real-world data.

In the world of generative AI and LLMs, Mistral AI is emerging as a strong competitor to OpenAI's ChatGPT product. Their most powerful B2B product is Mistral Large, which has already been sold to Microsoft for upward of $15 million to support their Azure platform. Their other product, Mistral 7B, an open source version of their foundational model, has also been used by companies like Hugging Face and is free to use. Companies like Mistral AI are taking the idea of powerful LLMs and customizing them for more specialized uses. This is a significant step in the direction of AI democratization because it allows for companies that want to adopt AI without having to front the significant cost, commitment, and time to develop their own LLMs. When companies as large as Microsoft accept this kind of partnership from third parties, it's clear that it's a solution that can help transform the B2B landscape when it comes to making generative AI more accessible.

Another great example of an AI company driving success in the B2B marketplace is a California-based gaming company called GGWP. While they don't create their own games, they use AI to reduce toxicity in the gaming culture and provide gaming companies with a dashboard where they can see how their users are performing from a moderation perspective. As gaming companies consistently take a more serious look at the safety and health of their gaming communities, they will come to rely on companies such as GGWP to ensure all their users feel safe. Moderation has long been a topic of discussion in social media companies and gaming companies alike, and it's really nice to find an AI company finding success in this field.

GGWP's most recent funding round in July 2023, led by SK Telecom Ventures and Samsung Ventures. won them $10 million in funding. Their success serves as a positive reminder of the importance of leveraging AI over human workers in specific fields. The biggest issue in employing human moderators is first and foremost the emotional toll it takes on their mental health as they go about their day. Scanning for hurtful or violent language isn't for the faint of heart, and finding depictions of graphic or hateful content on a consistent basis can deflate even the strongest of us. There's a strong case to be made for the ethical use of AI in helping us tackle the problems that we, as human beings, would rather not be tasked with.

There are also many examples of big tech companies that are investing heavily in expanding their B2B AI products. *Microsoft* has integrated AI into its business solutions in products like Dynamics 365 that provide enhanced **customer relationship management** (**CRM**) and **enterprise resource planning** (**ERP**) systems through things like leveraging predictive analytics, customer insights and behaviors, decision-making automation, and operational efficiency. Microsoft uses models like BERT and GPT (transformer models) for text understanding and language generation. *Google Cloud* has also been leveraging B2B AI offerings in healthcare applications by providing AI tools for imaging analysis, patient data management, and personalized treatment recommendations. In the case of Google, they use their own models (BERT and T5) for medical data extraction, summarization, and clinical decision-making, both of which are also transformer models. *IBM Watson* has also been expanding into financial services applications by leveraging AI for risk assessment, fraud detection, and customer service automation. IBM uses BERT, XLNet, and other tailored deep learning models for analyzing customer sentiment, interpreting contracts, and automating workflows.

When selling to other businesses in the B2B world, there's a standard operating procedure most businesses are comfortable with following. This includes everything from sales tactics to the approvals of budgets or statements of work to getting contracts approved through leadership, procurement, and all stakeholder teams in between. Because most B2B products are concerned with cost saving, productivity, and revenue generation, there are more structured interactions between all the players involved.

A couple of takeaways from the professionals are:

- Businesses looking to expand their B2B AI applications should start with core business needs first. Addressing specific business challenges, like improving CRM, operational efficiency, or risk management, is not only good business but will help you get traction from your business stakeholders. The closer you can align AI integrations and applications to business objectives, the easier the journey to getting buy-in will be.
- In many of these applications, AI is being leveraged to enhance human decision-making, so look for contexts where you can enhance predictive analytics and gather insights from large datasets that will help those within an organization come to a data-driven decision, particularly in cases where a business is helping to predict customer behavior or detect fraud. If we take predicting customer behavior further, we arrive at a deeper level of personalization. So, finding ways to apply AI functionality in ways that enhance customer satisfaction, loyalty, and lifetime value is also an elegant solution that is well worth the AI investment, particularly if you're able to leverage this at scale.

The spirit of a professional is competence and integrity, no matter what arises in the moment, and B2B markets can be thought of in many ways as well-oiled machines. Much like the products they represent, these companies themselves are looking for ways to optimize productivity, generate revenue, and lower costs in their sales cycles and market interactions. By starting with a clear strategy and selecting the right AI tools and applications, businesses can make a significant impact in their markets and elevate their products for the new competitive landscape.

The artists – examples of B2C products done right

The artists are here to show us how AI can be leveraged in a way that allows expression for potentially billions of consumers. B2C refers to products that are built with the expectation that the product will be bought and used by individual consumers rather than other businesses. With B2C products, there's a feeling that you're looking for a way to satisfy the needs and tastes of many individuals. Upon further reflection, B2C companies are looking to satisfy the few common collective needs of millions. The more due diligence that goes into imagining solutions and understanding unmet needs that would help so many people, the more prepared these companies will be to satisfy unmet needs.

Hands down our favorite AI-powered consumer app is TikTok. Considering the rise and prominence of the big tech giant's stronghold on American (and worldwide) consumers, the resounding success of its recommendation engine, and the controversial nature of the app's current standing with US lawmakers, we felt it was one of the more compelling products to include in this section.

Through a process called "frictionless embedding," their AI algorithm called "Monolith" is able to deduce preferences and interests based on real-time behavior. This allows their app to constantly be able to retain a myriad of varied and sometimes conflicting interests and offer its consumers videos that adhere to the evolving malleable and nebulous nature of their minds (to different degrees). One consumer, for instance, might love heavy metal music and through the app also discover that they also enjoy delicate, soft singer-songwriters as well. TikTok was initially thought of as an app that was popularized by younger generations for capturing their attention effectively, but today, it's much more than that. Consumers of all ages and backgrounds flock to the app to expand their horizons and explore their true interests without getting stuck in an echo chamber.

The Chinese giant uses the following types of AI to optimize its experience for its users:

- **Computer vision** to track images in the videos (YOLO, a convolutional neural network (CNN) model, and VideoMAE V2, a masked autoencoding framework with a vision transformer backbone, are used for pre-training video-related tasks like action recognition)
- **Natural language processing** to learn from sound and audio recordings with models like recurrent neural networks (RNNs) and Transformers being used to recognize and convert speech to text and CNNs being used to recognize and classify different types of sounds
- Metadata from the captions are used with **Transformer models** like BERT and GPT to best deliver content users find most compelling and reinforcement learning techniques to improve the recommendations on their FYP (for you page)

The biggest thing to remember with this example is it's a social media and entertainment platform, so we still see the inherent conflict between seeing compelling content and encouraging addictive behavior. This is still a problem area for the platform, judging by our own consumption. We're quite certain that we share this concern with our fellow one billion active users of the app, along with the company itself, which tries to limit the screen time of children under 18 to 60 minutes a day. As far as B2C products are concerned, we struggled to find a better or more successful example.

Now, let's turn our attention to another addictive app. This time, it encourages perhaps a healthier addiction: a thirst for knowledge! Pennsylvania-based Duolingo has leveraged deep learning in a way that helps people learn languages in a more efficient and specialized way for them. The gamified, personalized app reaches more than 300 million monthly active users and offers more than 30 languages. The main idea of Duolingo is it learns along with you and it offers you repetitions of words based on what it thinks you're just on the verge of forgetting. This repetition mimics the way we learn naturally and the gamified points-based behavior in the app gives you the emotional encouragement you need to keep going with your 15 minutes a day.

As a plus, Duolingo didn't start out as an AI company. It was founded in 2009, and once it started experimenting with personalization and A/B testing, it realized quickly that it ought to start testing with machine learning. Later on in the book, we will be discussing how to add AI features to an existing non-AI product. Duolingo offers us a successful example of a company that did just that. In particular, Duolingo uses RNNs for word sequencing and translating, LSTMs for recognizing long-term dependencies in text, CNNs for feature extraction and improving text understanding, Transformer models like BERT and GPT for understanding and generating natural language and providing contextual feedback, and reinforcement learning for personalizing learning experience and adapting content based on users' progress and behaviors.

Finally, one of the great resounding successes from the generative AI boom is an image generation app called Midjourney. In a rarity for the tech world, the company was able to generate $200 million in revenue with 40 employees and $0 in funding. Ever since the company's inception in 2021, David Holz, the founder of Midjourney, has been inundated with offers for funding from investors. It functions through the Discord chat app, which made ease of use and accessibility some of the hallmarks of its success. It's become a tool of choice for creative hobbyists and professionals who want to test out the capabilities of generative AI. The company's focus on sustainable, organic growth has also been a big part of its popularity, with subscription models between $10 and $120.

Another reason why we wanted to include Midjourney in this chapter is due to its inspiring labor practices. By having a democratic, flat organization and promoting a profit-sharing model among its admittedly very few employees, they're able to retain loyal employees by offering something many US companies cannot: true camaraderie. Even their disinterest in funding helps offer tangible and intangible rewards to the people who have built the company. With so many promising start-ups being snatched up by VCs, often with their own goals and priorities that conflict with those of the start-up they're funding, it's nice to see a company succeed when it goes against the grain. Often, start-up founders feel they won't be able to succeed without seeking funding, only to then see their company visions and missions compromised once they have it. In Midjourney, we see highly capable generative AI models being given to people, for a minimal cost, for the express purpose of delivering creative AI capabilities to artists.

We decided to call this group *the artists* because art captures the zeitgeist and varying energies coming from the social collective. It aligns pretty nicely with the idea of the invisible hand of the market. The concept of product-market fit is nebulous as well. The market wants what it wants and we all make up the market directly or indirectly, and the companies that create true art meet the market where it is. For instance, it's no wonder people are interested in learning languages.

The world is increasingly global. The global job market supports people all over the world who can work remotely or locally, so the idea that someone could move or just learn a language out of curiosity seems like it has increased within the last 10 to 15 years.

The idea that people would want to self-express in general, but particularly when sharing the collective retreat that was the COVID-19 pandemic, supports the idea of a global social media app that helps you create content almost effortlessly. TikTok and Duolingo captured the needs of the collective in a way that brought relief and ease to millions of people. With Midjourney, we also see an app that's focused on removing the fear-mongering of generative AI replacing artists by firmly and economically offering AI power to the very artists themselves. But even this is not without controversy. Companies like Midjourney and OpenAI are being sued by collectives of artists for copyright infringement because these models have been trained with real artists' work without their knowledge or consent.

B2C apps are in the business of tapping into the collective needs of the masses. Consumers are different from those working with other businesses. There isn't an entire village of potential gatekeepers getting in the way of finding an able and willing sale. When selling a product to consumers, you've just got to win over one consciousness. You'd be tempted to be all things to all people, in a way. But catering your product to potentially millions of viewpoints would be unsustainable. We can imagine it would also be hard to select and prioritize features. But capturing an emotional need that would inspire and engage millions of viewers would ring like a bell, and everyone who heard its echo would feel its radiance. How so very artsy!

The pioneers – examples of blue ocean products

When we think of pioneers, we think of those who go into uncharted territories with curiosity and are rewarded for their bravery. Blue oceans are especially difficult to navigate because you are creating demand for a product and, in essence, creating a market. The belief and vision required to survive the uncertainty in this kind of market is something very few people can sustain for very long. The evangelism, experimentation, and mission-building make pioneers stand out from the crowd. The first path to be forged carries with it an inherent virtue. All subsequent paths have you to thank for laying the foundation.

With pioneers, we wanted to focus on companies that were seeking new categories or use cases for their products. These are companies that are finding new paths to create demand for what they offer, and these companies create demand for themselves in multiple ways. This is what's referred to as a **blue ocean**: a competitive landscape that's still being formed, that doesn't see a lot of competition (yet), and that has to advocate for its own demand. It's relatively straightforward to explore the market in a red ocean because so many pathways have been created. There is already a thriving ecosystem to learn and navigate. But pioneers have to work for this knowledge.

The companies themselves have to do enough research and development to warrant the pursuit of this industry they're helping build. Blue oceans also put a lot of onus on the early players in their industries to create thought leadership and evangelism to keep the potential competitors, investors, and customers curious in the first place.

Bearing AI is a start-up that particularly caught our eye. Coming off the heels of a pandemic that has exacerbated the issues of an already strained supply chain crisis, we really loved the idea of a deep learning platform that was looking to not only streamline supply chain routes but also solve the problem of maritime fuel consumption. We still find it shocking that we are so reliant on physically moving goods such large distances, and that labor costs and raw material availability fuel this highly inefficient global trade system we've built. We hope that one day, we can create more efficient localized systems for managing global consumption, but in the meantime, there's Bearing AI.

What makes Bearing AI a good example of a blue ocean product? There really isn't much competition. A Google search for deep learning maritime route optimization products produces very few results, and this is the heart of the blue ocean landscape. You're not clamoring and infighting with your competitors for customers. The inherent challenge with blue oceans is that you have to create demand and justify the existence of your product. You have to essentially create a compelling story for why you should be in business, and why the world needs your product. It's a tricky game, but if you do it well, you get your pick of customers. And a compelling case Bearing AI did make, raising $7 million in August 2022 with a funding round led by AI Fund and Mitsui & Co., Ltd.

Initially, Bearing AI was in closed beta for a select group of customers, and recently, they were able to launch their new fleet deployment optimizer, which allows for predicting carbon intensity and emissions and minimizing vessels across a fleet. This is another advantage of the blue ocean: you can create an air of exclusivity. Only certain special people get to know about this new, exciting venture.

We've talked about the ethics of using deep learning models and the issue of explainability. This is something that becomes more or less important to the success of a product based on the use case. With Bearing, the deep learning models optimize for the best routes, and for the company and its customers, the accuracy of the models is more important than their explainability. If it works, it works. Here, we see less of a reliance on explaining how the models come to certain conclusions because their customers are more concerned with whether the routes make sense and are indeed more efficient and less concerned with how the inner mechanics arrive at certain conclusions. This might not be the case if it were, say, a dating site or a property tech company optimizing costs for renters.

Speaking of explainable AI, our second blue ocean example is Fiddler AI. We love that the rhetoric around explainability in machine learning has come to such a head that we're seeing companies emerge that partner with you to help create the transparency and trust needed to ensure that we continue to use AI in a responsible and ethical way. There aren't many companies doing this, and Fiddler AI allows you the ability to adhere to regulations and build models into your product at every stage of their life cycle. Fiddler AI does a great job of justifying itself: by bringing up the very real risks that might impact every company that's creating an environment for using machine learning models. They also do a great job of evangelizing their product. From the World Economic Forum, to Forbes, to CB Insights and Gartner, they're getting a lot of amazing press and they're using that to make a strong case for themselves in the blue ocean they find themselves in.

Part of the evangelism being done on the part of the pioneering companies also applies to the competitive landscape. When you're in a blue ocean, your competitors are your friends because they justify the existence of your product and solution. The more competitors you have, the more credibility you have as a business and the more you help create a supportive ecosystem with few key players. In many ways, pioneers are expanding and covering new ground in all directions.

The rebels – examples of red ocean products

Red oceans are markets that have become inhospitable environments with a full, mature, and developed competitive landscape. In this environment, many paths have already been formed and you have to choose one to live in and serve. You still have the option to create something new and you can employ many strategies to beat the competition and specialize, but you're challenged with having enough intuition to see what the next step should be. Part of the challenge with having such a diverse and thriving ecosystem of competitors is that it's hard to know what direction you want to really go in, or which competitors are truly taking a piece of your business.

Intercom, an Irish-founded company headquartered in San Francisco, has grown to become a major player in the customer communication and support software space globally. Operating in a highly competitive, red ocean environment, they face a crowded market of well-established platforms like Zendesk, Drift, HubSpot, and many others. These competitors offer similar customer engagement solutions like live chat, chatbots, AI-driven workflows, automated messaging, CRM integrations, and customer support ticket tracking. Feature parity, price sensitivity, and a focus on existing demand means that their strategies are limited when it comes to differentiating themselves in the red ocean they play in.

In an attempt to stand out, Intercom has highlighted their AI capabilities, such as chatbots and predictive tools, to streamline customer interactions and offer a more intuitive customer experience. They've also emerged as a customer-centric option due to their emphasis on a more human-centered, conversational approach to dealing with customers. Additional strengths, such as ease of use, a unified platform, a robust ecosystem of integrations and APIs, and community management, have also helped to secure their place in the market. Their significant investments in AI and automation have allowed them to automate routine tasks like answering common support questions and qualifying leads so that their customers and users can focus on troubleshooting more complex issues.

We see this with another prime example in a red ocean competitive landscape, Lilt, a machine translation company also based in California. The localization and translation industry has also seen its fair share of infighting; this is a market that's full of translation services and tech providers, and Lilt combines the two. From an AI perspective, what's special about Lilt is that it's *"the world's first and only interactive, self-learning neural machine translation system,"* according to their CEO, Spence Green. Its founders built Lilt as a response to the inadequacies they saw in machine translation while working at Google Translate, and their work has paid off. The company's most recent series C round raised over $55 million in funding in April 2022, bringing the total funding to more than $90 million at the time of writing.

The field of AI medical imaging has grown rapidly as the healthcare and biomedical fields have been huge areas of interest (and funding) for AI applications. One that stands apart from the rest is Gleamer AI. Their products AI Copilot, BoneView, ChestView, and BoneMetrics have been able to interpret, detect, and analyze images for a variety of findings. Hailing from France, their market penetration has been incredible, being trusted by over 1,200 health institutions, present in over 34 countries, and assisting with over 15 million patient exams in a year.

When we think of rebels, we think of courage. It takes a lot of courage to compete in a highly crowded marketplace. With so many competitors, you can easily get distracted and burnt out. Every competitor represents a potential direction in which you could grow your product. Other opportunities can be elusive if you're focused enough, and companies that try to be all things to all people suffer from this lack of focus. Product managers and leadership will need to work together to ensure there is a clear vision for the product and that the company can rally behind that vision as a team. To us, rebels are fearless, and exploring this fear could be helpful to understanding successful companies found in red oceans. Perhaps the fearlessness comes from the practiced courage of a company that's consistently facing so many adversaries in the market. Eventually, you stop sweating the small stuff and gain the gift of having laser-focused vision.

The GOATs – examples of differentiated disruptive and dominant strategy products

Now, we'll turn our attention to the main market strategies and see some of the *greatest of all time*, or *GOAT*, examples for each strategy. A market strategy informs your go-to-market team's efforts. Will you be going after customers that have too many options or not enough? Will you create a product that effectively works better or worse than your competitors? These seem like obvious questions when putting together a business plan, but once things get going for your company and you start getting some customers, suddenly these decisions might not be as concrete as when the company was first formed.

One of our greatest lessons from the start-up world was about getting comfortable with asking questions that seemed so baked into the company mission and ethos that they seemed obvious. We've asked these questions reluctantly in previous experiences but we don't anymore. Companies can change their strategies if they're tempted enough, but every change poses a potential risk to the company. Market strategy informs everything from how you build a product to how you sell it. A lack of clarity into these aspects of your market strategy could result in building a product that isn't right. A lack of clarity in the communication and sale of a product could result in acquiring customers that aren't right for the company you are.

The following chart from Wes Bush's book *Product Led Growth* offers us a look at the four chambers of strategies companies can use in their growth:

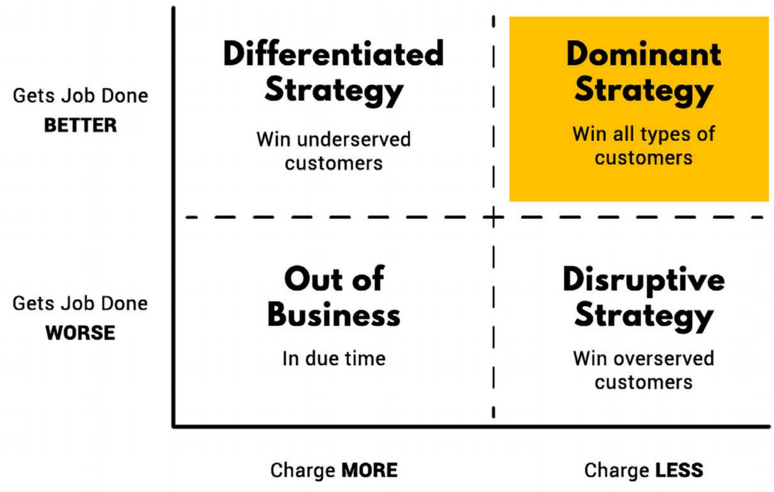

Figure 4.1 – The chambers of strategies

This chart shows us the strategy quadrant that delineates the various areas of a go-to-market strategy and highlights the key areas between the differentiated, dominant, and disruptive strategies. Every company will need to decide what its go-to-market strategy will be because this will inform how it sells, how it markets its products, and the language it will use to reach its target customers.

In the following sections, we will be focusing on the differentiated, disruptive, and dominant strategies, and choosing one example for each that's using AI to fuel these growth strategies.

The dominant strategy

In the **dominant strategy**, you are looking to win customers of all types: those that are looking for a superior product as well as those that are looking to pay less. This increases your market share overall and is why this strategy is referred to as a dominant strategy, because it's a winner-takes-all mentality.

Let's take the example of fast fashion, a highly competitive market if we've ever seen one. If you're a fast fashion retailer, it's not only important to have a cheaper product, but you must also have operations to support the delivery of that product more efficiently. Think Netflix, a good example in its heyday, but today we struggle to find a better example than SHEIN. The Chinese brand manifests its dominant strategy by leveraging AI to better anticipate new trends and predict demand for certain products in the market. It marries that data with its supply chain to ensure that it can deliver on changes in demand, to its customers' delight.

SHEIN uses autoencoders to enhance its recommendation system and optimize user preferences, transformer models like BERT and GPT for sentiment analysis, chatbot responses and information extraction from text, CNNs and YOLO for image recognition in product tagging and trend detection, LSTMs and RNNs for demand forecasting, K-means clustering algorithms for grouping fashion trends, reinforcement learning for dynamic pricing, and regression models for predicting price adjustments based on demand. This impacts everything from its marketing efforts to its in-app experience to the reviews of its products. The evidence is very telling as well.

Previously, a Spanish fast fashion company, Zara, had the fastest turnaround time of 3 weeks for creating and delivering a collection. SHEIN brought that down to a new best-in-class of 3 days – a truly admirable use of AI and machine learning. Though it may be effective, this example is also a complicated one because it highlights the nature of AI's effectiveness being in conflict with the world at large. Fast fashion already exacerbates environmental impacts and ethical concerns regarding the use of AI. Quick turnarounds and low costs lead to unsustainable practices, excessive waste, and poor labor conditions. But it's important to remember that while AI can improve operational efficiency, and it certainly has for SHEIN, it doesn't inherently address the fundamental problems fast fashion creates. In order to mitigate these concerns, it's crucial for companies to use AI in conjunction with ethical and sustainable practices.

Another company dominating a highly competitive ecosystem is Hugging Face, which not only brings machine learning code repositories, models, datasets, and web apps to demo AI-powered products to smaller companies, but it does that through an open source community. They've also recently partnered with AWS to offer their machine learning capabilities to their open source community as well, features that are otherwise only available to the AWS customers that can afford them. The proof of the feasibility of the company is in the funding; it raised over $235 million in a series D round in August 2023.

A dominant strategy is one that undercuts the competition in the market by getting the job done better and costing less. This is effective because it maximizes its rewards on both ends. It gets to receive the customers that want the job done better as well as the customers that just want to pay less. That's quite a lot of the pie you're winning. When done right, companies employing these strategies can essentially print money. But with great power comes great responsibility. There should always be a discussion about the tradeoff between ethics and financial success, and companies of all sizes need to remain vigilant about the impact they're having on the world.

The disruptive strategy

With a **disruptive strategy**, you're still selling a product for less money, but it offers you less too. Who wants that? People who are overserved and bombarded with options, but who actually need relatively simple tools compared to the competing options they have to choose from. We can see no better example of that than Canva. You can edit photos and create anything from a social media post to a resume using Canva, and the tools to do so are incredibly user-friendly and simple.

The Australian creativity platform leverages AI to offer customers more of the kinds of templates and content they're looking for. They do this by using CNNs for image processing tasks like background removal and object recognition, transformer models like BERT and GPT for auto-generating design suggestions and extracting text from images, GANs for creating new design elements and enhancing image quality, and reinforcement learning for optimizing design recommendations and layout suggestions. While it offers *less* than Adobe's Photoshop or the Microsoft Office suite, it does offer what its users are specifically looking for, and it's either free or cheap to use.

It quickly rose to unicorn status for expertly meeting its customers where they were. Canva did this by conducting extensive market research to understand the needs and pain points of potential users and analyzing existing design tools to identify gaps in the market and discover areas for improvement. It also segmented their audience based on design complexity, industry-specific needs, and whether they were professional or non-professional users, allowing the company to tailor offerings specific to different user groups. One of the biggest contributors to Canva's success was the freemium model, which attracted a wide range of users. Those who had more specific needs could address them with paid features, but the freemium model allowed Canva to meet users at different stages of their design journey.

What's most interesting about disruptive strategies is the ability to influence your potential customers by showing them a different angle. Disruption changes the paradigm of the existing power structures in a market and a new tastemaker arrives to inspire their competitors and customers alike to try new things with regard to how they ask for and use a product. This change in paradigm offers the entire market the gift of novelty and specialization. Maybe the job gets done worse, but this is beneficial for a new group of proposed customers. Perhaps this new group is one none of the other competitors thought to try and serve. Canva also offers us a glimpse into the potential danger disruptive strategies can face. Recently, the unicorn company was in the news for raising annual prices on their business-focused subscription service by 300% or more, driving users to threaten to abandon the software for competitors and claiming it's no longer the "simple and affordable alternative" that got them to buy it in the first place. Canva's response was that the new prices reflect the true value of their new AI-powered product experience. In this case, Canva appears to be abandoning its disruptive strategy for a more... differentiated strategy. This example is also a reminder that investing in AI is likely to incur costs that your company will then have to pass on to your users and customers. This decision may, in turn, affect your go-to-market strategy.

The differentiated strategy

In a **differentiated strategy**, you may sell a superior product that specializes in some niche, but you also charge more for it to account for this specialization. My differentiated strategy example is a company that hits close to home: a British machine learning-based property tech company I worked with named Beekin. Beekin isn't a property technology platform that offers cheaper services than its competitors, but at the time, it offered a next-generation platform no other property tech company was able to offer. We built a machine learning-native platform that did everything from making market evaluations to predicting future behaviors to offer an optimal price for renters. Our customers didn't have many options for competing tools because their alternatives were property tech giants with antiquated rule-based engines, and the competitive landscape in property tech is only just being disrupted by AI.

Differentiated strategies flourish in environments where they can best communicate the strength of their products. Because the price tag is also higher, there needs to be a compelling reason why someone would pursue differentiators in a market. Unique products that differ greatly from their competitors thrive in this market strategy. In the case of Beekin, there weren't many other competitors using machine learning as the foundation of their product, and it was able to help many people carry out their jobs. This is the greatest strength of a differentiated product: to satisfy a niche and advancing market. The need for differentiation is what furthers and matures industries. Once customers with more particular and pressing needs arise, differentiators pop up to show their competitors how they might begin specializing and rising to meet their customers where they are.

Let's take a moment to appreciate all the examples we have seen so far. The work and dedication that it takes to align a viable functioning product with a business model, a product strategy, and a loyal customer base is nothing short of a triumph, and success should be shouted out and celebrated. We learn so much from the successes of others, and we had a lot of fun choosing fitting examples for this chapter.

Whether you are trusting the brightest minds in AI or you're getting the help of a machine learning intern just out of bootcamp, know that a lot can be done with very simple models. Striving for algorithmic perfection and relying on simple tools that are commonly relied on and used are equally admirable. It's largely a question of preference and use case.

Summary

In this chapter, we covered some promising examples of AI done right in recent years for the purposes of finding inspiration in the successes and use cases of our peers in the space. As a product manager, having an understanding of your go-to-market strategy, business model, and the market you serve is crucial. We are firmly in the era of AI proliferation now. Many businesses are adopting AI within their products and their businesses, and we hope you've enjoyed these examples as we continue to understand the particulars of managing an AI product.

In the next chapter, we will put our futurist hats on and understand a bit about what the emerging trends on the horizon in AI are. We believe all product managers need to have a streak of futurism in their composition. Having a grasp of where things are in recent and present moments provides a snapshot with which to evaluate your existing strategy. But in tech, things evolve pretty quickly. In business, projections and future goals can change from week to week as well.

All of us will have to entertain future potentialities safely if we want to influence leadership, resources, and time toward the expansion of our AI products into areas that will inspire our customers and critics alike. Let's take a look and see what AI trends are in the forecast.

References

- *How Duolingo's AI learns what you need to learn*: `https://spectrum.ieee.org/duolingo`
- *Hazy company info*: `https://www.cbinsights.com/company/anon-ai\`
- *GGWP company info*: `https://www.crunchbase.com/organization/ggwp-65c2`
- *Intercom*: `https://www.intercom.com/`

- *How does TikTok use machine learning?*: `https://dev.to/mage_ai/how-does-tiktok-use-machine-learning-5b7i`
- *Canva*: `https://www.lifeatcanva.com/en`
- *Lilt company info*: `https://www.crunchbase.com/organization/lilt`
- *Hugging Face*: `https://huggingface.co/`
- *Mistral AI*: `https://mistral.ai/`
- *Gleamer AI*: `https://www.gleamer.ai/`
- *Midjourney*: `https://www.midjourney.com/home`
- *Canva says its AI features are worth the 300 percent price increase*: `https://www.theverge.com/2024/9/3/24234698/canva-price-increase-300-percent-ai-features`

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

5

AI Transformation and Its Impact on Product Management

As I look to the future and imagine what's to come regarding AI products, I feel tremendous inspiration and optimism about AI's ability to help us out of some of the greatest challenges that face humanity. I moderate a monthly talk with *Women in AI* called *WAITalks*. A lot of the discussions in these talks are centered on how the speakers broke into the field, what excites them most about AI, and what they wish others knew about their area of AI. There's a lot of beauty and insight in the different answers we receive from various women working in AI, from research to corporate to entrepreneurship.

One of our March 2022 speakers worked in AI policy at Meta. Because so much of her work involves ethics, we spent a great deal of time discussing the risks and inherent problems AI poses. But I also wanted to know what someone in her position looks forward to regarding AI, since it's always nice to get a sense of the opinions of diverse sets of people. Her answer blew me away – she said that because her work focuses so much on the legalities, risks, and biases AI poses, it would be easy for her to take on a pessimistic view of AI and its effect on humans across the world. Instead, she said she felt, more than ever, that AI will be foundational in helping us with the greatest problems humanity will face, from cancer to climate change to the social biases that plague us and keep us apart.

In this chapter, we will be discussing how we can help AI achieve its potential in helping us solve some of the biggest problems facing us. We'll explore a few waves of AI transformation and cover ideation, building **minimal viable products** (**MVPs**), and how to leverage AI for the common good as well as for commercial success. We will first take a high-level look at how AI will revolutionize our economic systems and how we relate to money and value. Then, we will take a look at the commercial level and how AI is fueling growth in MVPs. This will be an exploration into the kinds of integrations AI will enable at the goods and services level. We'll also discuss the expansion of AI in how we're governed, particularly the intersection of public citizen data and how governments will choose to use data on their own citizens. Finally, we will explore the use of AI in healthcare, and end with offering some insights into how AI is being used for the good of all and the fulfillment of basic needs worldwide.

With that, in this chapter, we will cover the following topics:

- Money and value – how AI could revolutionize our economic systems
- Sickness and health – the benefits of AI and nanotech across healthcare
- Goods and services – growth in commercial applications
- Government and autonomy – how AI will shape our borders and freedom
- Basic needs – AI for Good

Money and value – how AI could revolutionize our economic systems

AI has been no stranger to fintech as well as our greater economic systems. We've seen the rise of cryptocurrencies develop in concert with this change in how the financial industry operates. Smart contracts specifically have been a huge motivator for crypto and blockchain adoption. This isn't a book about crypto, but it does signal a change in how financial industries and economic systems fundamentally operate.

AI and **ML** are excellent tools for optimization, and there are few more compelling use cases for AI/ML than the optimization of profits. The financial industry has been leveraging quantitative analysts or *quants* for a long time. Quants do statistical and mathematical modeling for financial and risk modeling problems and, in many cases, it has been a profitable career for those who were interested in *the sexiest job of the 21st century*, aka data science, even before this term was made famous by *Harvard Business Review* in 2012. Modeling and risk minimization have been pushed into overdrive with AI and, considering the financial industry as a whole is people-heavy, the adoption of AI has been competing with human jobs.

Examples and use cases

Let's take the example of trading and the use of algorithmic trading. This is the ability to program a set of instructions to account for different variables, such as volume, price, and time of day. In comparison to a human trader, an AI or algorithm can optimize for the purchase of certain shares and learn from successful past behaviors much faster. Deep learning can also be leveraged here, and the output of its performance can be used as input to further train the algorithmic trading platforms, which means that these complex systems can start to see patterns and recognize opportunities almost infinitely better than a human trader.

There are several high-profile examples of this. Renaissance Technologies Medallion Fund uses AI/ML and quantitative models for trading and has achieved an average annual return of almost 40% since 1988. Goldman Sachs has integrated AI/ML into their Marcus AI platform, which provides consumers with data-driven financial advice and internal teams with AI algorithmic trading capabilities. BlackRock uses its AI platform Aladdin (Asset, Liability, and Debt and Derivative Investment Network) to manage its portfolios and optimize trading. JP Morgan developed its AI trading algorithm called LOXM to analyze historical trade data and optimize large block trades. The speed, data ingestion and capacity, adaptability to a changing market, and AI-driven precision inherent in these platforms have allowed them to maintain their competitive edge.

Let's look at some general use cases of AI that have been helpful in the financial sector:

- **Detecting and preventing fraud**: AI can help analyze transactional patterns to identify anomalies and flag suspicious behavior. Companies can use AI to detect fraud in credit card transactions, learn from historic transactions, and predict potential future threats to reduce fraud-related losses and improve customer trust.
- **Optimizing accounting and regulatory compliance**: AI can help automate financial statements, balance sheets, and various aspects of compliance processes, monitoring, and reporting. This reduces compliance costs, improves reporting accuracy, helps organizations adapt to regulatory changes, and minimizes operating costs.
- **Customer service and processing customer reviews:** Chatbots and virtual assistants can answer queries, provide account information, and process basic transactions, reducing operational costs and automating repetitive tasks. Feedback can also be tracked and used to improve operations and customer satisfaction.
- **Assessing risk and credit scoring:** AI can expand data points used to analyze risk and credit to include browsing behavior and e-commerce transactions to assess creditworthiness. This can expand access to credit and also provide more accurate credit risk assessments.
- **Automated loan underwriting:** AI can automate document review and decision-making based on predefined risk models. This can lead to faster loan approvals, less paperwork, and reduced bias, leading to more inclusive lending practices.
- **Personalized financial products:** AI can help offer personalized products based on customer needs, customer behavior, and financial data available to offer customized savings plans or credit offers based on customers' individual financial goals.

These are all areas of financial services that have seen AI transformation. In looking toward the future, we also see tremendous investment allocated specifically for AI expansion. A 2021 report by the **Organization for Economic Co-Operation and Development** (**OECD**) predicted that by 2024, global spending on AI would reach over $110 billion. In 2020, it was about half of this predicted amount, and given the impact of the COVID-19 pandemic, this trend toward AI adoption has been accelerated. According to the **International Data Corporation** (**IDC**): "*Presently, the global Artificial Intelligence market stands at nearly $235 billion, with projections indicating a rise to over $631 billion by 2028*". This means that we've already more than doubled the projections made in 2021 for 2024 and I would venture we're likely to see the same for 2028 figures.

Though there are often concerns about the *black-box* nature of deep learning models, we do want to note that in performance-driven environments such as trading, this isn't so relevant. The same goes for fraud detection and prevention, where the goal is minimizing fraud and financial loss. This is also true for portfolio management where portfolio performance is more important than methodology. However, assessing bias and making sure algorithms are not autonomously making decisions that would impact disparate groups of people will always be relevant, and we have governmental bodies that are looking into the fairness of financial exchanges and market regulation. These governmental bodies themselves are embracing AI as well – toward the end of improving regulatory activities, surveillance, and compliance more easily and faster.

Limitations and uneven adoption

There are still limitations to the use of AI/ML. These include incomplete, inconsistent, unreliable, and outdated data; model transparency is often a regulatory issue; bias and fairness issues stemming from the model's training data; ethical concerns; and the model overfitting to historical data in a way that doesn't generalize well for the present data. Modeling is currently being used in a capacity that still requires significant input from human staff.

In 2023, the U.S. **Securities and Exchange Commission** (**SEC**) proposed a series of AI rules and investigations into the development and use of AI models. While the SEC themselves are using AI to analyze complex data and assess market risk and potential fraud, they only use AI to help prioritize examinations for the **Division of Economic and Risk Analysis** (**DERA**); for more details on this, you can refer to `https://perkinscoie.com/insights/blog/artificial-intelligence-sec-proposals-and-concerns`. Despite these trends, and with the rise of more and more competitive LLMs crowding the market, we are still seeing AI stall when it comes to adoption in finance, with two notable exceptions:

- Asset management
- Mutual funds

These areas are the exceptions because the adoption of AI in fintech has been uneven, with these areas being favored. Asset management and mutual funds use quantitative strategies and are highly data-driven areas that are able to benefit from AI's ability to process and model complex financial data. Other areas of finance can face regulatory, security, and data privacy issues. Based on a recent *Harvard Business Review* article, the major disruptions that have come from a "*dominance of technology and data*" in this space have most favored the rise of "passive" fund managers rather than the historical trend of "active" fund managers. What this means is, over time, the influence of data and AI has led to managers trusting the growth and prosperity of indices rather than managers that charge hefty fees for "stock picking" based on "informational edges." This means that at a macro level, the gains that come from more data being analyzed at scale are surpassing those that come from more in-depth, micro-level analysis, as has traditionally been the case.

From financial transactions to the exchange of stocks, to boosting customer retention through personalization, to the regulation of financial activities and risk management, we've seen a few examples in this part of how AI is impacting the industry as a whole. With the rise of more accessible generative AI apps and pilots, we're going to see AI increasingly playing a vital role in compliance, customer experience, fraud detection, risk assessment, and financial analysis as they relate to financial sector activities.

Product perspective

We'd like to end with a more philosophical view of what AI adoption means for the very concept of value, worth, and money from a product perspective. If you work in fintech and you support financial products, you're likely concerned with not only the value this will deliver to your customers but also keeping an eye out for risks to the commercialization of your product. The financial examples in this section remind us of the importance of ideating use cases and products in a way that delivers value first and foremost.

Whether you work in Web3 and blockchain tech, traditional banks, or in the regulation of financial products, you're looking to harness AI/ML in a capacity where your product is helping someone improve some outcome. Help comes in all shapes and sizes. As you're in the throes of ideation and looking to offer something the market needs, consider products that are helping someone to arrive at conclusions in a faster, easier, or otherwise more efficient way. Maybe it's making your customers more profitable or minimizing their costs. Maybe it's helping your customers predict future uncertainty. Maybe it's offering them a way to expand their generational wealth and gain financial independence. Maybe your offering is using AI/ML to best understand your customers through personalization. There is immense value in this too.

If you're thinking about how to best leverage AI/ML for your fintech products, think in terms of how you can best deliver revenue and value to your end users and your business. Here are some approaches you can adopt:

- What degree of personalization can you offer your customers to build loyalty? Whether it's AI-powered advisors that offer investment or portfolio advice, tailored banking products and services, personalized loan offers, savings plans, and credit card recommendations, there's likely some way you can communicate value to customers in an otherwise dense and overwhelming market.
- How can you create a better customer experience that instills more trust from your customers? Customers can benefit greatly from agents and chatbots that handle customer queries, offer support, automate tasks, enhance customer service, proactively recommend solutions, offer market analysis, send targeted marketing campaigns, offer personalized rewards, loyalty programs, and predictive analytics, or analyze their behaviors and transaction patterns.
- Where can you make operational efficiencies for your business and others? Automating back-office operations, compliance checks, document processing, transaction reconciliation, data management, reporting, risk assessment, transaction monitoring, fraud prevention, and anomaly detection are all available options for AI integration.

Let's look at a different domain now, healthcare.

Sickness and health – the benefits of AI and nanotech across healthcare

AI continues to be incredibly successful when applied to the medical and healthcare industry in capacities like revolutionizing diagnostics, personalized treatment plans, and operational use cases. AI tools are being used to analyze medical imaging for X-rays, MRIs, and CAT scans at levels that surpass human productivity and accuracy. AI use cases for things like detecting tumors, bone fractures, and heart conditions have been successful because they're able to identify anomalies early, helping already short-staffed and stretched-thin medical systems in many countries.

We're also seeing personalization going far with AI. Tailoring treatment plans to genetic, environmental, and lifestyle factors for individuals has been successful because of AI's ability to analyze large datasets like genomic sequences. AI applications focused on oncology and cancer research are helping reduce the one-size-fits-all approach to cancer treatment options by optimizing treatment paths for individuals. Even the path to drug discovery and development is shortened by the use of AI applications that are focused on analyzing biological data to predict how new compounds will interact with the human body.

Even beyond the applied medical use cases, AI is still being used to optimize healthcare networks by helping to streamline administrative tasks like scheduling appointments, managing patient records and processing insurance claims, transcribing medical notes, and assisting with patient documentation, which is routinely done by doctors and other administrative medical staff.

Examples and use cases

Let's look at some use cases and examples that have been helpful in the healthcare sector:

- **Drug discovery and development**: Companies such as CytoReason, DeepCure, and BullFrog AI help their customers find and analyze new molecular compounds to create novel drugs for some of the most pressing illnesses that plague humanity, shorten time to market for these new drugs, and assist with the patenting process. Companies such as Standigm enable AI workflows to leverage AI through the journey of highly customizable target identification and lead generation. Before AI, the journey to discovering a new drug, conducting clinical trials, and bringing the drug to market might have taken upward of 10 years. Along the same lines, we're also seeing generative AI increasingly becoming a part of discovery trials to combat things like antibiotic resistance. Stanford Medicine came together with McMaster University to create SyntheMol, a new generative AI model that's responsible for creating chemical recipes and molecular structures to combat one of the leading pathogens that contribute most to antibacterial resistance-related deaths.
- **Personal health monitoring**: We also see personalization and prediction being applied to personal care through the use of AI, which is particularly exciting when you consider the doctor shortage we're experiencing worldwide. Wouldn't it be nice to see a personal health dashboard, where you can check your vitals on a regular basis through the use of nanotechnology and the analysis of data streams of your own personal data? Companies such as Biofourmis are investing in this and are also providing patients and doctors alike with the power to detect illnesses and potential health-related issues early, with their FDA-cleared Biovitals Analytics Engine. We can easily envision a scenario where companies start to build similar dashboards, customized for people with a specific illness or within specific demographic groups.
- **Hospital systems**: Even hospital systems are using generative AI models to help with emergency room visits and documentation during patient visits in high-traffic areas, as was the case with HCA Healthcare, Inc. in Nashville, Tennessee. As part of a pilot project with Google Cloud and Augmedix, a healthcare tech company, HCA plans to expand the time their physicians are able to spend with patients by limiting the aspects of their roles that are tedious and automatable. Another Tennessee health group based in Franklin, Community Health Systems, launched a partnership with a company called AvaSure to promote the use of virtual "sitters" that are available to patients deemed high risk and in danger of falling from bed rest.

- **Healthcare for disabilities**: Necessity is the mother of invention, and disability is also a huge area of healthcare that AI has been serving for some time. With the combination of robotics, computer vision, voice assistance, audio assistance, speech-to-text, and smart home tech, we have been applying and will continue to apply AI in contexts that can help individuals with disabilities related to sight, sound, smell, taste, touch, mobility, mental health, understanding, communication, and learning.

  A company called Parrots has created a medical assistive AI called Polly that's assisting people with neurological conditions. Polly helps its customers physically see the world around them better with 360-degree vision and navigate within their home or outside, offers real-time prediction by analyzing eye-tracking technology to anticipate needs, and facilitates the communication of needs, thoughts, and desires to the people around them. Essentially, it's a whole suite of products in one platform for a specific use case. Wheelmap is a crowdsourced navigation app that helps you find wheelchair-accessible routes, building on its knowledge through a community of users that participate in a feedback loop. Biped, a company that makes smart harnesses, has developed a product that fits like a backpack to assist blind and vision-impaired people with potential collision risks as they navigate the world.

- **Mental health and suicide prevention**: These are areas we've seen tremendous investment in with the use of AI applications. I teamed up with a colleague from my data science program to create my dream product in 2020. It was a mental health-focused conversational AI that got to know you, established goals with you, and recorded your mood. We were interested in using AI to serve as a real-time journal and dashboard for self-reflection, with the option to share this data and dashboard with a therapist if you chose to work with one.

  Through the process of building the idea and MVP, we did our research and discovered other companies that were tackling mental health and suicide prevention in their own ways. Initially, we wanted to focus on all consumers but saw that the language models would have to be tailored considerably to the needs of more specific groups if they were going to be effective. Companies such as ResolveX are using AI to work with suicide hotlines to eliminate manual and administrative tasks so that the workers can focus on the human on the other end of the line. Woebot developed a chatbot using cognitive behavioral therapy techniques to help people battling depression and anxiety.

- **Vision enhancement**: Google's DeepMind achieved notoriety for its ability to help detect retinal diseases and cardiovascular risk factors by using AI on retinal images. A company called Be My Eyes partnered with OpenAI to incorporate its GPT-4 model into their Virtual Volunteer product, an AI-powered visual assistant for those who are blind or have low vision. Since its inception in 2015, Be My Eyes has connected its staggering 6.3 million volunteers to assist users with daily tasks but now that they've incorporated a foundational model or LLM into their product, they're able to escalate toward their goal of improving "accessibility, usability and access to information globally" (`https://www.bemyeyes.com/blog/introducing-be-my-eyes-virtual-volunteer`).

- **Wireless implants:** In February 2024, a Neuralink chip was successfully implanted into the brain of a 29-year-old paralyzed man named Nolan Arbaugh, making it the first wireless AI product to be implanted into a human subject. This milestone brings accessibility to those suffering from debilitating health conditions that limit physical mobility. When asked to describe how it works, he likened it to using "the force" from the movie Star Wars on a cursor (`https://www.wsj.com/tech/neuralink-shows-first-patient-using-its-brain-implant-device-67a8b03a`).

Product perspective

In healthcare, we see a similar mission-driven element to the ideation, MVP, and release cycle, but the good of all is once again the main focus because, as is often the case, these products can have significant repercussions for users and customers if the products themselves haven't been thoroughly tested, and if the downstream effects of using these products haven't been studied. While it's always important to get the right idea and a working version of your idea out there to test with users, in healthcare use cases, it's imperative that all areas are strongly accounted for because you're dealing with people's health and safety.

Based on the examples discussed in this section, between diagnostics, discovery, research, and quality of care, there are many areas where AI can help us improve how we administer health and well-being in our world. As we saw with the global pandemic of 2020, and continue to see in these subsequent years, the healthcare systems and applications that serve us today are increasingly inefficient and fragile in many countries, independent of their development status. If you're thinking about how to best leverage AI/ML for your healthcare products, think in terms of how you can bolster performance, shorten time to care for patients, and improve patient outcomes. Here are some approaches you can adopt as you're conceptualizing products in this space:

- What degree of personalization can you incorporate into AI use cases that will help improve patient outcomes and patient experiences? Personalization is being used currently to optimize individualized treatments by analyzing patient data at the individual and collective levels. Can you help patients identify new treatment options or plans, particularly for complex conditions like cancer or rare diseases? Focus on ways your AI application can help with forecasting future health risks or help with preventative care.
- How can you better improve long healthcare processes like drug discovery and research? Reducing R&D costs and shortening the time to market for new drugs and treatments is a good way to bring products to market that have high ROI and adoption outcomes for customers, users, and patients.
- Could there be ways you can help operational efficiency across the medical system? Between insurance companies, providers, clinics, and hospital systems, there's enough bureaucratic red tape to go around and often it's admin folks, doctors, and patients themselves who are left with the brunt of the work of getting their coverage needs and costs met. If your product can reduce the burden on any of these major players, the market will surely reward you.

We will continue to see more healthcare-focused AI applications in the coming years, particularly now that generative AI models have become so much more accessible. While these developments are exciting in their own right, efforts need to be made to ensure the benefits of AI are reaped by organizations and patients equitably. Ensuring AI benefits are properly distributed and contribute to the accessibility of care will not only spread the wealth to those who need it most but will also help those who are AI skeptical to trust that the tech is serving them and not the other way around.

Once you have a few good ideas, you're ready to test. You can build MVPs that test out the effectiveness of those ideas in a way that communicates value and aligns the collective benefits of the target market you're serving, bringing commercial success. We will explore this in more detail in the following section.

Goods and services – growth in commercial applications

Now to the fun stuff! One of the great blessings of AI/ML is seeing just how creative we humans can get with it. In the previous section, we wanted to start with finance and business because this is the heart of the capitalist free market, and product management is an inherently commercial role. Now let's take one step away from that and briefly explore the world of AI commercialization for consumer goods and services. We will use a few key examples across various industries and sectors that we find particularly inspiring as promising, early MVPs.

In keeping with the theme of this chapter, we encourage all AI PMs to keep up with promising developments in AI/ML products because this phase of AI adoption we're in is all about inspiration. This is the golden age of the entrepreneur, and nowhere is that more evident than in the start-up world. Despite short-term hurdles like interest rate accessibility, bearish investment and speculation, and macro trends that lead to companies conducting mass layoffs, we maintain that this period of time is especially rich for AI implementation, particularly now that generative models have become so widespread. Technologists from all ages would relish the opportunity to be around in this era of big data and AI. As we've mentioned previously, the earlier phase of AI was research-heavy and we were starved of data. Now that we have loads of data and a pretty good understanding of the models and languages we can leverage to make the most of it, we can actually apply it and do what humans do best – create!

I had the privilege of participating in the **Women in AI Accelerator** (**WaiACCELERATE**) in 2021, and after seeing all the candidates going through the program and the applications going through the process, from ideation to building an MVP, I found myself inspired by the creativity going into these nascent ideas. From optimizing building materials to scanning the internet for fake news, to robotics, to AI assistance advocating for women, **BIPOC** (**Black, Indigenous, and People of Color**), and other marginalized groups, it seems like the sky's the limit for AI adoption. If you have the curiosity or opportunity to join a hackathon or an accelerator of this nature, whether as a participant, an organizer, or an advisor, I highly recommend you take it. These environments are incredible for expanding your creative horizons with AI and they're often in need of people with diverse backgrounds who can offer novel ideas and approaches to AI.

Examples and use cases

Now let's look at the following examples across some different domains:

- **Ideating for sustainability**: One company that's leveraging AI in innovative and creative ways is Symrise with their Philyra product, an AI-powered perfumery made in conjunction with IBM. Philyra's library consists of 3.5 million legacy formulas, over 2,000 raw materials, and 20 dimensions to help it come up with new recipes to suggest to its customers. Symrise's approach to AI was also an interesting one. Instead of opting to embed AI for improving efficiency or automation, a buzzword that means very little in applied terms today, they wanted to take a more inspired philosophy when it comes to AI and embed it firmly into the creativity of their perfumery while still honoring their human perfumers, not myopically replacing them. They're also using Philyra to bolster sustainability by having it optimize for ingredients that have "improved renewability and biodegradability credentials" (`https://www.symrise.com/scent-and-care/competence-platforms/philyra/`).
- **Personalized AI companions**: Replika is a chatbot that gets at the deepest part of what makes us relate, an AI companion and friend who "cares". Obvious data privacy concerns aside, Replika serves as an interesting companion because it learns from your interactions to better understand and anticipate your emotional needs. You can chat about how your day is going, do activities together, and have a soundboard to bounce all your ideas off of. Since the first edition of this book, Replika has expanded to more use cases with its product, allowing for video calls and even shared **augmented reality** (**AR**) experiences, bringing Replika into the real world. Considering the loneliness epidemic among both men and women, and the social and emotional vulnerability of people across ages and generations, Replika offers people something that doesn't come easily with the complex and unpredictable humans that are in their lives: to be understood and fully accepted. As generative models get better and better, they'll be changing how we interact with all kinds of tech, simply because they're making it easier for us to interact with it. And because having an AI assistant to use when you're making or working on virtually anything makes your job easier, you're more likely to adopt it.
- **Creating art**: Now, let's move onto a topic we think of as being uniquely human – making art. It's a worthwhile philosophical topic to ponder. Can we consider something art when it doesn't come from the tortured soul of an artist? I think we will come to find, increasingly in the future as well, that indeed we can. AI-generated art is a controversial topic and it's one that's received a lot of scrutiny from artists, as we covered in *Chapter 4*. Copyright issues aside, there are merits to the accessibility of creative AI. They can offer brainstorming assistance to artists and be used as tools. But even beyond that, the poignancy of a poem isn't erased just because the source, or even inspiration, isn't human. This point brings up the importance of finding ways for AI to coexist with human artists. Will we see AI art rise to such prevalence that humans are left out in the cold? Or will human artists start to incorporate AI into the art they themselves create? There has already been pushback on the part of consumers and artists alike who are not only hesitant to use AI but are morally against it. Though their intentions are in the right place, it's likely the artists that can use AI in a way that doesn't threaten their success who will pave the way forward.

> There's a whole slew of music apps using AI to generate new music, such as SOUNDRAW, OpenAI's Jukebox, and Boomy, or lyrics, such as Audoir. Or, maybe we can explore the myriad ways AI has been changing how we optimize film schedules, promote films, predict the success of films, edit films, and even produce films. The use of AI-generated content was a big part of the bargaining done by the Writers Guild of America (WAG) and SAG-AFTRA during the strikes in the US in 2023. For months on end, countless people came together to protest the networks that were looking to diminish the labor and creativity of the actors and writers that make blockbusters happen. There is virtually no creative industry that's going to be untouched by AI in the future, and this trend has only escalated with the rise and prominence of generative AI models. The rise of LLMs and other generative AI models has not only contributed to a wealth of accessible creativity, but they've severely sped up the ability of AI to effectively convince you. What these powerful models offer us as technologists, AI PMs, and customers is the ability to interact with a machine that doesn't feel like a machine. The use of these new models and abilities will become more prevalent in the future as AI applications become more social. We surpassed the threshold of a previous limitation of generative AI models to generate long-form text, which is huge.

By now, generative models have been around long enough that we've all seen AI-generated selfies, images, videos, and songs. To the novel eye, they may be fleetingly entertaining and fascinating. But after you've seen a few, you've seen them all. No doubt the quality of the outputs will improve as the models get better. But I maintain it's much more interesting to see an artist use AI like they would any other medium to not only bring forth art into the world (and get properly paid for that art) than it is to spin up a tired, formulaic AI-generated movie in the likeness of Brad Pitt.

All the products we've mentioned in this section started with an idea to create a product that could bring the richness of life to consumers, and these product teams built MVPs out of these seminal ideas.

Product perspective

An MVP should be released to market quickly so that you can start the process of collecting feedback on how your users are resonating and working with it, allowing you to iterate further through the production process. In many of the examples that we just discussed, companies built MVPs that they then tested, and they found an overwhelming willingness and curiosity from their target markets to the point where there are now many competing products for each aforementioned area. Building an MVP that you can test and improve quickly not only helps your own business but also creates an even bigger market that attracts more competitors.

To best leverage AI/ML to build an MVP, here are some approaches you can adopt:

- Begin with a narrow use case of AI and align that use case to the needs of your business
- Leverage pre-built AI models where you can from open-source libraries or AI platforms like TensorFlow, Hugging Face, or Google Cloud AI
- Agree on data requirements for the MVP and, once it's live, continuously collect data on how users are interacting with it to fine-tune the product experience and keep improving

- Leverage customer feedback regularly to inform your product development cycle once the MVP is live
- Agree on a baseline of explainability for how your MVP is arriving at determinations or conclusions that you will pass along to your users and customers
- Leverage natural language where you can to help with adoption. If people have the option to use natural language and effectively speak to a piece of technology in their own native language, they'll gladly take it.

Building an MVP is a tool to give you an early signal that you've got the right idea, market, and target customer base. This helps you to quickly and successfully balance the needs of your customers with your own commercial success. We can also call this product market fit. However, how does this change when applied to governmental sectors? We will explore that in the following section.

Government and autonomy – how AI will shape our borders and freedom

Using AI is not all fun and games. The immense availability of data and AI simultaneously opens us up to the risk of authoritarianism and the rise of oppression, and the collective empowerment to combat threats to democracy across the world. At the moment, we have some striking contradictions of power and protection being exerted in very different ways, based on a system of shared societal values. Governments will increasingly need to focus on preserving their country's ecosystems and the fallout from climate change, and with that focus, there will need to be some degree of investment as well. The private sector can only do so much, and we hope that governmental bodies will use their influence for the betterment of their citizens. Let's look at how AI will be leveraged in governmental capacities and what it means for all of us as consumers and makers alike.

Let's take an issue such as *voting*, which, at least in the US, has seen many headlines during the last few presidential elections. Companies such as Microsoft are developing electronic voting systems, such as ElectionGuard, as part of their Defending Democracy Program to allow for safe, secure, and verifiable voting processes in the future.

AI also has a role to play in *policy-making* itself. There are several stages in which AI and big data can be leveraged to facilitate the process of legislating. We see a range of use cases where AI can be helpful at the legislation and policy levels:

- **Pattern detection and identification of topics that constituents most want to address:**
    - The UK government has a Data Science Campus that uses AI to analyze social media and other sources to detect patterns and identify key issues the public is most concerned about. This has allowed the government to understand public sentiment about a variety of topics and assist with policy-making (`https://datasciencecampus.ons.gov.uk/`).
    - Singapore's government launched the "Smart Nation" initiative to use AI to analyze data from various sources, including feedback directly from citizens to detect emerging trends and topics of public concern to target public engagement strategies that are relevant (`https://www.smartnation.gov.sg/`).

- **Creation of analysis and forecasting at a granular level:**
    - Estonia's e-Residency program uses AI to help non-Estonians start and manage a business online in order to attract more e-residents and investors from out of the country (`https://www.e-resident.gov.ee/`).
    - Australia's "National Map" analyzes geographic and socioeconomic data to help facilitate better urban planning, resource allocation, and disaster management by offering real-time insights for initiatives and responses (`https://nationalmap.gov.au/`).
- **Generation of data to support findings:**
    - The U.S. Census Bureau uses AI to process and analyze large volumes of data from the census and surveys it administers in order to provide accurate population estimates and demographic data that supports policy decisions and resource allocation (`https://www.census.gov/data/data-tools.html`).
    - Canada's "Open Government" data initiative helps process, generate, and make data available from various public sources to enhance transparency and encourage public participation in government (`https://open.canada.ca/en`).
- **Implementation and diagnostic feedback for how certain policies are performing:**
    - The Netherlands "AI for Government" project uses AI to monitor and evaluate the implementation of policies and provide diagnostic feedback on their effectiveness, helping to improve the impact of initiatives and legislation (`https://nlaic.com/en/about-nl-aic/`).
    - South Korea's "Digital Government" project uses AI to track and analyze the performance of various policies and programs and help increase public satisfaction with policy outcomes (`https://www.dgovkorea.go.kr/`).
- **Evaluation of the ultimate effectiveness of certain policies:**
    - New Zealand's "Policy Project" uses AI to assess the impact of various policies by offering simulations and predictive modeling to help refine policies and enhance economic outcomes (`https://www.dpmc.govt.nz/our-programmes/policy-project`).

If governments choose to partner with the private sector or invest in their own development of AI applications to address policy-making, we will likely see an expansion of the democratic process.

On a broader societal level, we can see some novel applications of AI that will likely continue into the future. The alleviation of bias is something AI can help address at the governmental level, where policies are impacting varying demographic groups and the intersectionality of their experiences. Companies such as Algorithm Audit are doing just that with their bias detection tool. Their product works by first analyzing the data used to train and test a model. It also uses various statistical methods to assess and visualize different types of bias to examine whether certain groups are unfairly represented or treated by the model. Finally, it then offers recommendations and remediation efforts to mitigate the bias and monitor the models over time.

Part of the ethical conundrum of using AI resides with the issue of bias in algorithms, a topic made famous recently by M.I.T. Media Lab researcher Joy Buolamwini in the Netflix documentary *Coded Bias*. Joy has taken her work forward by building her own organization (and effectively her own movement), the Algorithmic Justice League, to continue this advocacy for algorithmic bias further. We're also starting to see companies such as Pipeline Equity that advocate for pay transparency and equity across demographics to minimize the pay gap and give organizations sound business proof for why they should invest in pay equity as well. They do this effectively by developing AI tools that use data to detect and analyze bias in hiring, **diversity, equity, and inclusion** (**DEI**), promotions, and pay equity. They also learn from the data they amass by developing guidelines, best practices, research, and standards for regulatory frameworks and building training resources to spread awareness and educate about risks and biases in AI.

By far, where governmental bodies are concerned, AI is the elephant in the room. Particularly, the use of autonomous weapons and mass surveillance are most controversial for understandable reasons. People want to feel safe, but they also don't want their privacy abused. I had the honor of speaking with Colonel Dave Barnes, chief AI ethics officer at the US Army Artificial Intelligence Task Force and professor and deputy head of English and philosophy at the USMA, at a virtual conference I participated in with Women in Trusted AI in 2020. The greatest insight from that conversation was the acknowledgment that every country is still grappling with how to most responsibly and effectively use AI and understand the potential harms and risks. It's clear that there's much to be done across governments to set limits and understand the full capacity of how AI is leveraged. We're seeing more conversations about how to limit the negative externalities of AI in *defense*. Regarding autonomous weapons, the US Department of Defense's principles of AI are as follows: responsible, equitable, traceable, reliable, and governable. This sentiment is perhaps best expressed by Air Force Lt. Gen. Jack Shanahan, director of the Joint Artificial Intelligence Center: "*The complexity and the speed of warfare will change as we build an AI-ready force of the future. We owe it to the American people and our men and women in uniform to adopt AI ethics principles that reflect our nation's values of a free and open society*." Societies aren't free or open if they're being controlled and oppressed.

The same could be said for mass surveillance. AI can be used to empower governments to oppress and control their citizens, and it can also be used to keep them safe. Mass surveillance is also bleeding into the private sector, with many private companies choosing to use facial recognition, which is very difficult to opt out of, for something as simple as going to a theater or a bar. The recent **AI Global Surveillance** (**AIGS**) Index shows an increase in AI surveillance being used worldwide.

The following is an extract from *The Global Expansion of AI Surveillance* by Steven Feldstein (`https://carnegieendowment.org/research/2019/09/the-global-expansion-of-ai-surveillance?lang=en`):

> *"The index shows that 51 percent of advanced democracies deploy AI surveillance systems. In contrast, 37 percent of closed autocratic states, 41 percent of electoral autocratic/competitive autocratic states, and 41 percent of electoral democracies/illiberal democracies deploy AI surveillance technology. Liberal democratic governments are aggressively using AI tools to police borders, apprehend potential criminals, monitor citizens for bad behavior, and pull out suspected terrorists from crowds. This doesn't necessarily mean that democracies are using this technology unlawfully. The most important factor determining whether governments will exploit this technology for repressive purposes is the quality of their governance – is there an existing pattern of human rights violations? Are there strong rule of law traditions and independent institutions of accountability? That should provide a measure of reassurance for citizens residing in democratic states."*

With the rise of generative AI models, governments will have to invest in *data synthesis and evaluation* regulations to ensure the responsibility of monitoring data that comes out of LLMs and other forms of generative AI. While these models are especially good at generating new data, making sure that generated data is correct, safe, and legal is another story. Increasingly, we're seeing issues with LLM providers being sued for violating intellectual property laws, and this is a trend that's likely to continue and impact artists and creators whose content has been used to train LLMs and whose content is likely to be included in outputs shared with users as well. Companies like CloudFactory are promoting "Human-in-the-loop AI solutions" to combat the lack of accountability when it comes to long-term, scalable AI systems.

What does this mean for AI PMs? Well, for one, if you're in the *authoritarian tech* space, business is booming, no matter your politics! More importantly, it means that we as AI PMs have a responsibility to support and create products we believe in. Our input and involvement in building these products has massive repercussions, depending on the audience and whether you have your eyes set on governmental or military aspirations. You should think critically about the products you bring to market to see how you can combine the common good with your own success. These governmental examples also went through the process of ideating, from building an MVP to releasing it and getting feedback. The only difference is these products are being administered in areas that affect public citizens, and we would say the responsible use and development of these products is the hardest area to get right when the potential impact is so high.

Basic needs – AI for Good

Having started my career in sales and account management, I took its competitiveness for granted. But one thing that really struck me about the AI/ML community when I first joined was how collaborative and benevolent it was. There were so many open-source projects you could get involved in. You could put your projects up on Kaggle and win competitions. You could find an almost infinite supply of solutions to problems on Stack Overflow. I found a lot of success reaching out to people who shared their perspectives on their area of AI. Never was this more apparent than when I reached out to potential speakers to host a panel or a workshop to teach others. It makes this field all the more compelling because if you do want to focus on a particular AI solution, you can find lots of friends along the way to join your cause and support you.

Saving ourselves from ourselves is big business for AI, and there are many companies that put humanitarianism in their business model, using a combined passion and entrepreneurship to bring products to market that also benefit society on some level. You'll likely see the term *AI for Good* used by organizations that look to use AI for some social or humanitarian cause. Many of them work with other companies, schools, governments, and NGOs collaboratively to address certain causes. The **AI for Good** nonprofit organization has worked on smart cities, workforce diversity, AI ethics, and Economists for Ukraine projects alongside the United Nations **sustainable development goals** (**SDGs**) SDGs. The UN itself has the **International Telecommunication Union** (**ITU**), a specialized agency for information and communication technologies that has its own AI for Good year-round platform to host volunteers and serve their communities as well. There are more ecosystems, such as Omdena, DataKind, and Teamcore, that look to collaboratively address humanitarian causes as well. Microsoft, H2O.ai, Google, and IBM also have their own AI for Good initiatives.

AI for Good is such a wide area covering everything from environmental protection and disaster, water and food scarcity, and disasters caused by nature or war, offering assistance and accessibility for a variety of impairments, sustainable development, resource management, public safety, and many more. Let's look at some examples from different domains:

- **Water systems**: Nonprofits like Earth05 are creating an open-source community of scientists, academics, innovators, start-ups, and organizations that can use AI to predict and analyze future water systems and consumption, optimize water delivery systems, and create better irrigation systems toward the goal of positively impacting over 500 million people by 2030.
- **Agriculture and food distribution**: Sustainable food distribution and urban farming projects are being powered by AI, where companies such as Square Roots and Bowery Farming can make sustainable indoor farming more popular and widely adopted, particularly in cities where access to local fresh produce is challenging and continuously in demand. Companies such as Peat are able to help farmers in insecure areas better maximize their yields with their product Plantix.

- **Disaster relief**: This is another area where we see quite a lot of investment when it comes to AI for Good use cases, particularly in situations where there are search and rescue needs, harsh weather conditions, fires, and other natural disasters. Autonomous drones like Fotokite's Fotokite Sigma and robots like Shark Robotics' Colossus will become increasingly relevant as we experience extreme weather patterns exacerbated by climate change that put human first responders at greater risk.

Bringing products to market that affect people's access to basic needs is one of the greatest expressions of AI products we see out there, and we want to remind you of the importance of divergent thinking here with regard to ideation. Where AI for Good is concerned, we are barely scratching the surface of what AI can do, and we would recommend people who enter this realm as AI PMs focus most of their attention on the ideation phase. Once you've found a worthy application of AI for some social good, practice the use of divergent thinking. Hone in on the end result you're looking to achieve with AI, and use this as your guide in your ideation phase. Building an MVP will come with its own considerations, and you're already crafting this offering for the benefit of all and combining it with a way to create commercial success with your tool, so you can really take your time with ideation.

Summary

This chapter concludes *Part 1* of this book. We started with an introduction to AI and the infrastructure required to support it, went into the weeds of model maintenance and the particulars of ML and deep learning, saw a mix of applications and business model examples of AI products, and concluded with a glimpse into where AI is going next. We saw how AI can transform products by enhancing decision making, automating processes, and personalizing user experiences. To achieve this, PMs need to think about how their products best create value and how AI can maximize that value. While AI can leverage large datasets to demonstrate the effectiveness of products and policies, PMs must be aware of how bias and ethical considerations should be handled in their products to avoid controversy. Products across sectors are also susceptible to an evolving landscape of regulatory requirements. This was by far not an exhaustive list in terms of the future of AI, but I believe it's representative enough of where we're heading to offer those in AI product management a few insights into a few major sectors that shape our world through AI. AI is often discussed as an integral part of the Fourth Industrial Revolution, but when we consider just how many facets of our world and our businesses AI will impact, it's easy to see why we are constantly returning to the gravity of these terms. We know that we will continue to see more novel ideas of how to harness AI, and that's what makes this period of AI implementation so exciting.

Part 2 will expand on AI-native products by focusing on what it takes to understand, ideate, create, and productize AI. We will focus on practical applications and case studies that explore how AI products can be customized and how performance can be optimized, as well as going into some examples of common pitfalls and successes an AI PM can run into with the AI-native product. In the next chapter, we will be looking at what areas are essential when you're building an AI-native product. We will be looking at the particulars of managing a product from the ground up, with certain AI considerations and what AI PMs will need to account for as they begin the process of ideating, building an MVP, and ultimately, launching an AI product. Though we touched on some of these concepts in this chapter already, we will be expanding on them in the next chapter primarily for those who are entrepreneurial and are looking specifically to ship products that will meet the needs of their market using AI.

Additional resources

If you're in the US and at risk of self-harm, please check out the following resources:

- Crisis Text Line: Text *CRISIS* to 741741 for free, confidential crisis counseling
- The National Suicide Prevention Lifeline: 1-800-273-8255
- The Trevor Project: 1-866-488-7386

For those outside the US, the International Association for Suicide Prevention lists a number of suicide hotlines by country. You can find them by going to their website (`https://findahelpline.com/i/iasp`). Also, check out Befrienders Worldwide (`https://www.befrienders.org/need-to-talk`).

References

- *The role of big data, machine learning, and AI in assessing risks*: A regulatory perspective: `https://www.sec.gov/news/speech/bauguess-big-data-ai`
- *Artificial intelligence, machine learning, and big data in finance*: `https://www.oecd.org/finance/financial-markets/Artificial-intelligence-machine-learning-big-data-in-finance.pdf`
- *Data scientist: The sexiest job of the 21st century*: `https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century`
- *Bowery Farming*: `https://bowery.co/`
- *Square Roots*: `https://www.squarerootsgrow.com/`
- *Plantix*: `https://plantix.net/en/`
- *Peat*: `https://peat.technology/en/`
- *IBM AI for Good*: `https://www.restack.io/p/ai-for-social-good-answer-ibm-ai-for-good-cat-ai`
- *Google AI for Social Good*: `https://ai.google/responsibility/social-good/`
- *H2O.ai AI 4 Good*: `https://h2o.ai/company/ai-4-good/`
- *Microsoft AI for Good*: `https://www.microsoft.com/en-us/ai/ai-for-good`
- *Harvard University Teamcore*: `https://teamcore.seas.harvard.edu/`
- *Omdena*: `https://omdena.com/`
- *DataKind*: `https://www.datakind.org/`
- *AI for Good*: `https://aiforgood.itu.int/about-ai-for-good/`
- *Woebot Health*: `https://woebothealth.com/`
- *Resolvex*: `https://resolvex.com/`
- *Biped unveils an AI copilot for blind and visually impaired people at CES 2022*: `https://www.designboom.com/technology/biped-ai-copilot-blind-visually-impaired-people-ces-2022-01-10-2022/`
- *Wheelmap.org*: `https://wheelmap.org/?locale=en-us`
- *Parrots*: `https://www.flyparrots.com/`

- *Biofourmis*: `https://biofourmis.com/`
- *Standigm*: `https://www.standigm.com/about/company`
- *Bullfrog AI*: `https://bullfrogai.com/`
- *Deepcure AI*: `https://deepcure.com/our-approach`
- *Cytoreason*: `https://www.cytoreason.com/company/`
- *The global* expansion of AI surveillance: `https://carnegieendowment.org/files/WP-Feldstein-AISurveillance_final1.pdf`
- *DOD adopts 5 principles of artificial intelligence ethics*: `https://carnegie-production-assets.s3.amazonaws.com/static/files/WP-Feldstein-AISurveillance_final1.pdf`
- *Algorithmic Justice League*: `https://www.ajl.org/`
- *Coded Bias on Netflix*: `https://www.netflix.com/title/81328723`
- *Algorithm Audit's bias detection tool*: `https://algorithmaudit.eu/technical-tools/bdt/`
- *New cyberthreats require new ways to protect democracy*: `https://blogs.microsoft.com/on-the-issues/2019/07/17/new-cyberthreats-require-new-ways-to-protect-democracy/`
- *How artificial intelligence is used in the film industry*: `https://smartclick.ai/articles/how-artificial-intelligence-is-used-in-the-film-industry/`
- *AI Lyrics and Poetry Generator*: `https://www.audoir.com/`
- *Boomy*: `https://boomy.com/`
- *OpenAI Jukebox*: `https://openai.com/blog/jukebox/`
- *SOUNDRAW*: `https://soundraw.io/`
- *OpenAI updates GPT-4o, reclaiming its crown for best AI model*: `https://www.zdnet.com/article/openai-updates-gpt-4o-reclaiming-its-crown-for-best-ai-model/`
- *Breaking new fragrance ground with artificial intelligence (AI): IBM Research and Symrise are working together*: `https://www.symrise.com/newsroom/article/breaking-new-fragrance-ground-with-artificial-intelligence-ai-ibm-research-and-symrise-are-workin/`
- *How AI-enabled Chef Watson is bringing creativity and innovation To cooking*: `https://bernardmarr.com/how-ai-enabled-chef-watson-is-bringing-creativity-and-innovation-to-cooking/`
- *Renaissance Technologies: The highest-performing investment firm in history*: `https://quartr.com/insights/company-research/renaissance-technologies-and-the-medallion-fund`
- *IDC's worldwide AI and generative AI spending – Industry outlook*: `https://blogs.idc.com/2024/08/21/idcs-worldwide-ai-and-generative-ai-spending-industry-outlook/`
- *What the finance industry tells us about the future of AI*: `https://hbr.org/2023/08/what-the-finance-industry-tells-us-about-the-future-of-ai`
- *HCA Healthcare collaborates with Google Cloud to bring generative AI to hospitals*: `https://investor.hcahealthcare.com/news/news-details/2023/HCA-Healthcare-Collaborates-With-Google-Cloud-to-Bring-Generative-AI-to-Hospitals/default.aspx`
- *Computer vision for brain disorders based primarily on ocular responses*: `https://pmc.ncbi.nlm.nih.gov/articles/PMC8096911/`

Part 2

Building an AI-Native Product

Understanding what it takes to manage an AI program is a prerequisite of AI product management. Now that we've covered the basics of that in *Part 1*, we can move on to contextualize the output of that AI program when building products natively with AI.

This part consists of six chapters addressing the various relevant areas when going to market with a new product built with AI from the very beginning. We will cover the basics of creating an AI-native product, productizing the AI/ML service that powers it, positioning that product for various groups, AI/ML options you can take when building, and how that all relates to performance benchmarking, costs, and growth. We'll understand both the tactical and strategic considerations of designing, launching and managing an AI-native product by the end of this part. We will contextualize all these concepts with the help of a case study that will span all the chapters in this part.

This part comprises the following chapters:

6

Understanding the AI-Native Product

In this chapter, we will go over the essential components for creating a strategy for building an AI product. This strategy will allow companies a process that will help them succeed in building an AI-native tool from the first **minimum viable product** (**MVP**). This chapter is primarily for **product managers** (**PMs**), technologists, and entrepreneurs coming into the AI space who want to build or manage AI products that are native to AI. In other words, products that are built with AI from the start. In *Chapter 2*, we briefly introduced the new product development stages, and we will now build on that structure by focusing on the most important phases of introducing a new AI/ML native product: ideation, data management, research and development, and deployment.

We will also address the main contributors to your AI/ML product team, as well as the tech stack that will empower them. While data is foundational for AI-native products, the roles you fill to support the team responsible for their creation will be critical to their success. When building an AI-native tool, it's important that you do your due diligence so that you aren't being wasteful of your company's finances and resources. A bad hire or incorrect tech investment can be costly in both time and resources. By the end of this chapter, you'll have a firm understanding of the most important considerations for every stage of your AI product development lifecycle, as well as the development of your product team and the tech stack that will support them. We'll also discuss some of the differences between traditional product management and AI product management, as well as some of the relevant elements when customizing your product for different audiences and how to craft a message that resonates with them. To show you how this translates to the real world, we'll introduce a case study that we will continue to engage with over the course of the next few chapters.

We will be covering the following sections:

- Stages of AI product development
- AI/ML product dream team
- Investing in your tech stack (further expanding on concepts from the preceding infrastructure section)

- Productizing AI-powered outputs – how AI product management is different
- Customization for groups – considerations for verticals and customer groupings
- Selling AI—product management as a higher octave of sales

Stages of AI product development

Whether your organization is robust enough that you're supporting multiple product teams for various stages of your AI product development or lean enough that you're running on a skeleton crew that will see your product through each stage, you'll want to be cognizant of what these different stages are so that you can define success through each phase.

There are various schools of thought on what product management is or should be. Few universities offer courses on product management as part of an undergrad or graduate program. Since the product world is still nascent and growing every day, I would even venture to say that the role of a PM varies from company to company in terms of expectations, roles, and responsibilities. We'll do our best here to describe the various phases in a way that best summarizes the core aspects of AI product development.

Either way, as an AI PM or leader of an AI product, you'll want to consider how your product relates to each of these phases so that you can identify the phase your product is currently in and what you need to do to bring it to full maturity. Here's a tailored version of the key phases in AI product development:

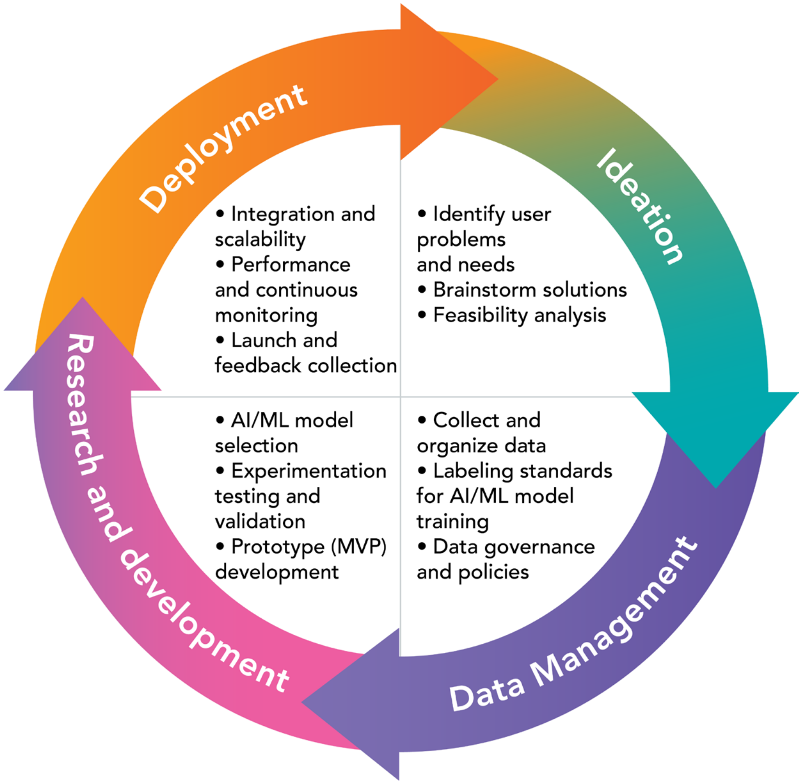

Figure 6.1: Stages of AI product development

Phase 1 – Ideation

Just as with the traditional software product lifecycle, as a PM, you'll be involved in the ideation phase. You might see this referred to as the innovation or design phase as well. Whatever you call it, this is the phase where the brainstorming on features is happening, and as a PM, you're a key stakeholder in influencing the direction of the product. This is the phase where you're also weighing the cost/benefit of incorporating various features, as well as identifying the non-negotiables for your product.

One differentiator between traditional software development and AI product development in this phase is that the focus on ideation is even more important before continuing because the costs of investing in an AI product are great. Therefore, you want to be heavily focused on this phase before you move on in the product development lifecycle. You don't want to be going back to the drawing board after you're in the thick of it with your product only to realize halfway through that you have a better idea of how to leverage AI for your product and have to lose months of work, time integrating a specific tech stack, storage, and an expensive headcount.

Ideally, the work in the ideation phase results in getting a picture of what the product will look like using UX mockups/wireframes and surveys, as well as gathering the requirements you need to have in order to properly scope out your MVP. As a PM, perhaps the most important part of your role is understanding what the actual opportunity or problem is you're addressing. This is true from the market perspective as well as the ML perspective.

Where ML is concerned, you're using AI to address the following types of problems:

- **Anomaly detection**: Applying AI/ML to identify rare or unusual patterns that do not conform to expected behavioral patterns (e.g., fraud detection)
- **Clustering/peer grouping**: Applying AI/ML to group data points into distinct sets based on similarities (e.g., user segmentation)
- **Classification**: Applying AI/ML to assign labels to data from an established set of categories (e.g., spam detection)
- **Regression**: Applying AI/ML to predict a numeric value based on training data (e.g., predicting housing prices)
- **Recommendations**: Applying AI/ML to suggest relevant items or content based on established preferences (e.g., shopping cart recommendations)
- **Ranking**: Ordering or prioritizing items based on criteria (e.g., search engine results)
- **Optimization**: Finding the best solution from a set of potential solutions (e.g., driving route optimization)
- **Generation**: Creating new data/examples based on input examples/prompts (e.g., image generators)

Getting clear on which Problem with a capital-P you're addressing in the market, as well as which ML problem you're solving based on the eight competencies we just discussed in the preceding list, will be your most important task as a PM. Everything else is secondary.

You'll also want to bring everyone together at this phase. Your UX designers, data architects, data engineers, data analysts, data scientists, ML engineers, backend/frontend/full stack engineers, and leadership team will make up the key stakeholders that you'll want to be involved in this phase so that every voice has a chance to be heard.

By the end of the phase, you should have the following decided:

- **What is the problem/opportunity you're addressing?**

  Identify the core challenge or opportunity you're trying to solve. Is there a gap in the market or is it just a pain point you're improving upon? Understand the underlying cause and the potential impact you're trying to have because this will set the foundation for your product.

- **What are the necessary outcomes/requirements to have a working MVP?**

  Defining a minimum set of features that will solve the core problem while meeting user needs will give you the blueprint for an MVP. Determine the essential functionalities for this first version, along with user flows and technical requirements. Create an understanding of what success looks like for this first version and remember that you have an entire product roadmap where you can make improvements and expansions for future versions.

- **How will this MVP be delivered?**

  Establish a delivery plan that includes team structure, timelines, resources, tools, technologies, and platforms that will be involved. Make sure business constraints align with your ability to deliver technically.

- **Who will use this MVP?**

  Create a map of your target users and personas to clarify their needs and pain points. Focus on your early adopters since these will be the people using this MVP. Might their needs be different from those of later users? This will be important information for you to have as you build future iterations.

- **How will they use this MVP?**

  Create a map of scenarios and usage patterns to understand how groups of users will interact with the product. Those use cases will help contextualize how users will interact with the product and what considerations are most important for them.

- **How will success be measured?**

  Establish KPIs and metrics early to evaluate how your MVP is performing with your early adopters. How will you be able to know the MVP is addressing the needs of your users and is ultimately solving the problem? These metrics will set the foundation of your feedback loop for future iterations of your product.

Phase 2 – Data management

Once you have a clear idea of what you're building and why, you invest in setting up your product for success. Throughout this book, we refer to AI, ML, and **deep learning** (**DL**) products, but it's important to remember the original colloquialism for any of these products is this: **data product**. At its core, any product that leverages AI/ML in any way is a data product first, and no data product can be good at anything without lots of quality, clean data. Organizing and expanding on the infrastructure needed to support all your data is the first applied step post-ideation. You'll want to understand the requirements, constraints, and vocabulary required for a system that supports the data gathering and data processing that will be required to make your AI product function properly because you'll be central in helping make decisions about this.

In this phase, you're also defining the best features to use in your AI/ML models. We discussed feature engineering in previous chapters. As a refresher, this is the idea that you're looking for facets or *features* in your dataset that you're primarily going to use for training your model. You can think of features as the individual columns of data elements that will be most important to include in your model preparation and training. You won't want to leave this up to your data scientists, engineers, and architects either. As a PM, you'll want to be intimately involved with and understand the various features that will be selected as the most relevant and final inputs in training the models that will make up your product.

Once that's been decided, this is the phase where the data pipelines to support this data feed will be put in place. Having a good grasp on how you're setting your organization and product up for sound data collection and preparation, data storage, and data pipeline practices will be crucial in this step as well. These data practices will allow you to leverage data for internal use cases and traditional analytics, and they will also serve as the true foundation of your AI product. While there will often be a data team handling these considerations, it is important for AI PMs to have an understanding of how these decisions impact the efficacy and cost of their products from a resource management and planning perspective.

By the end of this phase, you should be able to answer the following questions:

- **What kinds of data do we need to make this MVP successful?**

  Identify the types of data required for training and model development to make sure the models you eventually choose will be set up for high performance and accuracy.

- **Are there infrastructure needs that your organization doesn't currently invest in to support the success of this MVP?**

  Ensure data pipelines and storage solutions at the organization are sufficient to handle the volume and variety of data your MVP will need to be properly supported. Which constraints or limitations affect your data infrastructure?

- **What data quality standards will your MVP need?**

  Having a plan to maintain quality standards on the data you're using to train, validate, and test your models will be important to get the performance you're looking for from your MVP downstream. Make sure processes are in place to measure and maintain data quality throughout the lifecycle of the MVP as well as future iterations of the product.

- **How do different data decisions impact the overall cost and resource management for this MVP?**

  Document the various costs for decisions impacting data collection, processing, pipelines, and storage that will affect the budget and time for the creation of this MVP.

- **Are there legal or ethical requirements for the data you're gathering and using?**

  Work with your legal team to ensure you're maintaining compliance and regulatory requirements like GDPR and CCPA.

Phase 3 – Research and development

So, you've got your plan, you've got your data organized, and you're ready to start building. In this phase, you're researching and developing the actual structure and substance that will make up your AI product. **Research and development** (**R&D**) can be synonymous with experimentation. You aren't going to build a product around one model, and no matter how intuitive and talented your data scientists or ML engineers are, they likely won't go with the first model that occurs to them. We've gone over the various ML models in previous chapters, and we've mentioned a few times that most of the time, ensembles of models are what get used in products.

In the *Deployment strategies* section of *Chapter 1*, we introduced the concept of **A/B testing**, which is sometimes referred to as split testing. Essentially, the process is relevant here as well because there will be a fair amount of A/B testing and evaluation in this process, so you'll want to understand the basics of A/B testing as well as data distributions, cohorts, confidence intervals, and other probabilistic concepts.

This phase also means there's a fair amount of experimentation that needs to take place before a model is selected for your specific use case. If you're hiring the most experienced data scientists and ML engineers and you see them going back and forth with different models, seemingly unable to choose and wanting to further experiment until they see performance they can get on board with, this is normal.

But hiring a talented, knowledgeable, and experienced headcount is only one side of the coin with AI products. Sure, they have the know-how, but they aren't business experts. They're experts in modeling and creating algorithms that will fit the use case the business experts outline for them. Use the R&D phase as a way to give your data scientists and ML engineers the proper direction, and manage expectations with them. You'll want to properly set these hires up with the proper tools, stakeholders, and resources to make them successful. There are other considerations as well. Let's take the notion of performance. As a PM, it could be up to you, rather than the data scientists, to decide what level of performance or precision would be most acceptable for the customers you're trying to serve. Establishing these thresholds and expectations is something you, as the AI PM, will be regularly doing with your team.

By the end of this phase, you should have the following in place:

- **Documentation of the methods and models that have been tested**: This should include strong model contenders that have been evaluated along with metrics used to assess their performance. Criteria used to select the final model or ensemble of models should explicitly define the business and technical factors that led to the selection. Include content about how these models have been tested and if there are concerns that they may not scale or generalize well in real-world conditions. This should also include documentation of any A/B testing results, cohort analysis, and confidence intervals used for various models.
- **Clear acceptable performance thresholds for your MVP:** Include thresholds for precision, accuracy, and other performance metrics that are acceptable to internal stakeholders and end users alike. This will help make sure the MVP is performing well after deployment, and integrating triggers, logs, and alerts when performance dips below this threshold will ensure your team remains proactive should performance issues arise. Identify potential risks such as bias or performance degradation over time.
- **Tools, resources, or infrastructure that data scientists and ML engineers need to make the MVP successful:** In addition to the infrastructure needed to support your data management, there will also be infrastructure needed to support the AI/ML teams. Ensure your technical team has what it needs to carry out effective R&D including platforms, compute power, and other tools to support their work.
- **Product requirements and user stories that tie model selection and performance back to business goals and customer needs:** Often, when we're building and thinking about technical investments and restraints, we are looking to optimize our resources. But remember to connect the tech back to what matters most: providing value to your target users. Create a narrative around the research being done and how it's impacting customers, users, and the overall business's success.
- **A plan for refining and iterating on the MVP:** Remember the MVP is just the starting point. From there, you are making future iterations on your product. How will the MVP continuously improve? How will your models be retrained? How often? New data and performance feedback will start to create momentum around your MVP, so have a plan in place to catch this momentum.

Phase 4 – Deployment

We're using deployment as an umbrella term here for everything that comes after the R&D phase. At this point, you've settled on a model or a collection of models that satisfies your product ideation and MVP requirements, is supported by your data infrastructure, and has been thoroughly tested. Now, you're ready to build the supporting infrastructure to make sure this model and its outputs can be integrated into the broader product that your end users will experience. This is the deployment phase, where you're marrying the work of your data scientists and ML engineers to the work of your full stack developers and actually integrating it into the greater product that will support these AI/ML outputs.

Your product may have many *features*, and only a couple of them may be *AI features*, or the very heart of your product could be ML. Regardless of how rooted in AI your product is, you will need to deploy the ML findings from the models you've invested time in building in the context of the greater UI/UX. Because of this integrative dance that has to happen between your AI/ML code and your production code, you might go through your own R&D phase with deployment as well. Selecting the proper ML model and technique is a separate process from selecting the proper way to showcase and use the outputs of that ML R&D and how to maintain those outputs.

The broader context of the deployment phase is really to create continuity and delivery of the substance of your ML outputs. This is also where you'll be creating processes and structure around the continuous maintenance and delivery of your product, which were outlined in *Chapter 2*.

Here are some aspects that you should have taken care of during this phase:

- **Considerations for integrating your MVP into your product or platform:** Ensure the frontend and backend systems are fully integrated into your MVP so that AI functionality and UI/UX are blended seamlessly. Document the deployment process and details about how models were integrated into your MVP so that knowledge is retained and easily transferable.
- **Production infrastructure is set up for model serving:** Models used in production need to remain scalable, reliable, and responsive. Maintain the appropriate microservices, APIs, containers, or platforms needed to keep model deployment smooth in production.
- **Automation of model monitoring and maintenance:** Set up triggers, logs, and flags to monitor your models in production and alert you when performance drags below the thresholds you established. Tracking performance and detecting potential issues can trigger model retraining or updates when they're needed most. Make sure your MVP is set up with pipelines that ensure the model powering it stays relevant as new data is collected and used. Maintain performance standards by stress testing models to ensure that production performance metrics align with expectations set during Phase 3.
- **Feedback loops from customers and users post-production:** Whether it's in product surveys, questionnaires, customer calls, or focus groups, having a method of gathering and applying user feedback on your MVP will help identify areas for improvement and minimize edge cases or unintended bias.

Now that we've covered the various phases of the AI product development lifecycle, let's focus on the folks that will make the magic happen in the first place: the AI product team. Building a new team for a new AI product is tricky, and it can be tough to adequately hire for the needs of your product. Every product and company will have different needs, so we will try to remain as objective as possible in this section.

AI/ML product dream team

In this section, we will spend some time understanding the various roles that will empower your AI product team to maximize success. We'll also group these functions across the stages we outlined in the preceding sections so that we can get a sense of when these roles come into play. Note that not all these roles will be necessary in your organization. Every organization is different and will have different needs. Most companies operate under constraints and it's rare that we work with a fully staffed dream team, but this section is helpful in understanding which roles are critical to build an MVP. Use your discretion when building your AI teams. You may include other stakeholders in your AI program, but the following is a relatively complete list of the main stakeholders you will want to include in your hiring process as your AI/ML product team grows.

We'll now look at a cumulative list of roles that will likely apply to your ideal AI dream team. We have listed the roles in order of common appearance.

AI PM

Here, we start on our journey of creating a dream team. All organizations are different, and here are some approaches they may adopt:

- Some won't hire PMs until they're further along with building an MVP.
- Some might hire *founding* PMs.
- Some might already be mature organizations and are just ready to launch their first AI product.

No matter where your organization lies, you will eventually need all the activities required to support an AI/ML product to be centralized into one person.

A product is a living, breathing ecosystem in and of itself, so having a point person to go to who can be responsive to and responsible for the needs of various parts of that ecosystem will be important. If you're an entrepreneur, technologist, or investor who's hiring for this role, finding someone who's worked with or specialized in AI/ML products in the past will serve as your due diligence for making sure your product is properly set up for success. We consider it ideal to hire for this role before you start hiring for other roles to maintain a cohesive vision across your product. Having an AI/ML PM on board before you begin the implementation of your product means that person is aware of the historical discussions and market influences to support your product best.

If you're a PM interested in focusing on AI products, build your technical aptitude in AI/ML while improving your ability to make decisions using data. This will be a role where you will manage diverse teams of people, directly or otherwise, and your cross-functional leadership skills will be tested. Develop business acumen and AI ethics awareness to make sure you are bringing products with real value to market.

AI/ML/data strategists/architects

This will be more of a consultative role for most companies. Ideally, this person would help with the acquisition of talent, along with developing methodologies and workflows that would support your AI/ML function/team in its entirety. This might be a role that you hire for at the start of your AI/ML journey to help with key decisions about who to hire for various roles in your AI program or product team, what should be included in your tech stack, and how certain workflows should be set up with regard to experimentation and deployment.

This may be someone who stays on for the planning and initial phases of your product development (first 6–8 months) or someone who you keep on staff as a PM in more resourced environments. This person will also ideally be well-versed in AI ethics principles. If you're a start-up and you're creating a company and a product from scratch, this might be your technical cofounder who is able to act as your technical decision-maker. This role could also be referred to as a data architect role.

Data engineer

We have an architect for how the AI product/program team will function; now, we start laying down the foundations for our data pipelines with a data engineer based on the blueprint that's come out of our collaboration with an AI strategist. If you're setting up an AI product from scratch, look for a data engineer who can help support your team's choices of data engineering philosophies. Note that this person should not be looking to the AI PM for the proper data setup. Rather, they will inform the AI PM about what various data engineering choices will mean for the AI MVP.

Will you be using a data mesh, a data fabric, a data lake, or a data warehouse? As you're starting this function from the beginning for an AI/ML native product, you'll want to hire someone with the confidence and experience to help guide you with a proper setup. Migrating and changing your data architecture can be expensive and time-consuming. If you get the help you need properly at the beginning, it will save you potential headaches later and give you a sustainable, scalable infrastructure to grow with confidence.

Data analyst

A path for many companies will be to build on an already established foundation of data analytics before they start expanding on those analytics with ML. The traditional order of operations is gaining clarity and control over your data by hiring amazing data analysts first who will quickly figure out what's interesting and worth exploring in the data you already have. In this role, you're looking for someone who's able to quickly analyze your swathes of data with curiosity and exploration.

This is the person who might be your expert query maker, working in concert with your data engineer to feed queries through the workflows you establish to power your AI/ML product. These will also be the hires that will improve on the analytics you use internally to improve internal processes as well as the analytics you pass on to your customers through your product or platform's dashboard UI/UX.

Data scientist

Once you have the ability to generate and move data, you're ready to start experimenting with that data, and you'll eventually need to collaborate with someone who's an expert at model building. In order to be an expert at model building, you have to also be an expert at ML algorithms, big data, statistics, and programming. Most companies are coding in Python, so unless you're serving academic circles that primarily code in R, we recommend sticking with Python. This role doesn't just need technical skills—it also needs soft skills because it's a collaborative role, and it also communicates heavily with leadership.

At this point, you're hiring a data scientist to execute the business goals you've established early on. This person should have strong business acumen and must be able to understand the impact and purpose of their work. Though this will be a highly collaborative role, particularly early on during the AI MVP, they should have the expertise to technically and functionally offer the best models and solutions that work best with your organization's data and product goals. This role is inherently elusive to fill because finding someone with competence and experience in all these areas proves to be difficult for most organizations, particularly when you're in the early phases of creating a start-up or a product from the beginning.

ML engineer

At this point, you've got a firm handle on your strategy, your data architecture, the capabilities of your customer/product data, and the models you'll be using in your product. Your goals and objectives have been outlined and confirmed by the technical stakeholders on your leadership team. Taking a cumulative approach, by this point, you can outsource the execution of this planning-heavy part of establishing your AI/ML product or program. ML engineers are able to use the models and algorithms whipped up by your data scientists by incorporating (coding) them into the workflows or code repository for your product.

Again, because this is a role that will support the AI-native product, this role will have the burden of getting in the weeds with your data and algorithms to see what comes out. This person can expect to do a lot of trial and error as they feel their way toward acceptable performance. Your ML engineer will use the data that's been vetted by your data analysts and marry it with the models your data scientists have selected in their work. Nothing really matters until you actually deploy something into production. The act of putting it all together and ultimately deploying the code that supports all the prior functions falls on the applied ML engineer who makes the dream team. Depending on your organization, you might have to make a choice between a data scientist and an ML engineer.

While both data scientists and ML engineers know how to train, test, and select models for discussion with their stakeholders and collaborators, ML engineers will be the ones best equipped to be able to deploy them into production. The reason for this is that there are higher technical expectations for ML engineers than data scientists.

Frontend/backend/full stack engineer

Various manifestations of the preceding five roles exist in many capacities. Your particular organization might not need them all, all at once. But one thing's for sure: no matter how you navigate the previous five roles, you will always need full stack engineering to support the AI/ML work that needs to be done before it's time for deployment. ML engineers are able to add the AI/ML functions to your code base, but you'll need built-out engineering teams to support your product end to end. Your frontend, backend, and full stack engineers will collaborate with your ML engineers and data scientists often for the AI/ML native product.

Structuring these teams so that they're working intimately and have deep trust built between them will be important. Most of the time, any AI/ML product will have plenty of features and needs that don't have anything to do with the AI/ML aspect of your product. You'll need a built-out engineering team that manages the tasks and epics in your sprints outside of the AI/ML function. These are the folks that are going to get you set up with a working MVP of your product that you can iterate on as you continue in your product management journey. You don't have a working product until you're in this phase.

QA/testing engineer

One of the recurring themes of this book is about the iterative process of AI/ML products and the importance of maintaining performance metrics across AI products, and this point is perhaps best summarized with the QA or testing engineer role. These folks are the ones making sure the AI models and systems that support them are functioning correctly and are meeting the needs of the product. They are regularly testing the outputs of AI/ML models against performance standards and are validating that the data pipelines are properly feeding those models regularly.

They work collaboratively with data scientists and engineers to ensure the AI system that powers your MVP is running smoothly and they flag bugs, performance bottlenecks, and potential issues when it isn't. Depending on their AI/ML skill set, they may also support issues with model drift, data bias, and model performance issues over time. If there is also automation you're introducing into your AI system, these are the folks who will implement those triggers, alerts, or flags to catch issues faster.

UX designer/researcher

The manifestation of your MVP is your product's starting point. You can't iterate a product that doesn't yet exist, and reaching MVP level allows you a confident start. Depending on your organization and philosophy, you may already have leveraged UI/UX feedback during the planning phases. Perhaps you did have a product designer on board early on, or perhaps you worked with a UX researcher when you were doing your early market reconnaissance. But once you have a working MVP, you're able to establish a baseline with your users (or beta users) and get a sense of how they're receiving value from your product.

UI/UX is all about how to most efficiently deliver value to your end users. Minimizing sources of frustration and inspiring moments of delight is the primary focus of this function. A good UI/UX designer or researcher will guide the visible changes of your product that your end users will experience. This role is a critical partner to an AI PM because it can help validate customer pain points or market strategy decisions. While the go-to-market and product teams might be making decisions based on data, UX research can offer that objective third-party perspective. We recommend having this function be an ongoing engagement so that you can learn from your users as you build further iterations of your product that align most closely with what your users—and, by extension, their buyers—will approve.

Customer success specialist

Having proper channels for listening via UI/UX design and research is important, but eventually, your customers will want you to talk back—and not just through the product itself. You might not have a huge need for a customer success team if you're creating a B2C product, but if you're selling an AI/ML product to other businesses, you certainly will. The customer success team exists to make sure your customers are properly using your product and receiving the intended value from it.

They're also a great source of feedback for your data and AI/ML functions. A customer success specialist offers the feedback loop that an AI/ML product will need to build and improve over time. They're also incredibly helpful at helping to expand on product features and collaborating with products on which potential new features end users might really love to see. If you're an AI/ML PM, make sure you have a warm and open relationship with the customer success team because they're your biggest source of knowledge on the ground. They're the ones actually talking to your customers!

Marketing/sales/go-to-market team

This is another one that will likely be involved at the beginning and throughout your product development journey. We're lumping in sales and marketing as **go-to-market** (**GTM**) because this is essentially what this function does. It communicates out to your market: the entire landscape of your potential customers. This is where you'll combine your skill sets to optimize your product language fit or the message your market will see about your product. You'll be using your GTM meetings to decide on the following:

- Terminology about your product
- How you're going to position it in the market
- How much of your proprietary tech and algorithmic magic you're going to allow your sales force to discuss with potential customers on demos
- How much explainability you're going to offer
- How you're going to communicate about technical decisions you've made about your product
- How much information about your product you're going to shield from the world in the spirit of competition

As you build and improve your product, your GTM team will stick with you, and you'll be regularly in the flow of discussing how you're going to communicate about upcoming features or releases.

Before we move on, let's summarize all the critical and non-critical roles that support the successful launch of an AI MVP:

Role	**Description**	**Department**
	CRITICAL	
AI Product Manager (AI PM)	Defines the product vision, requirements, and success metrics. Aligns AI features with business goals and user needs.	Product Management
Data Scientist	Designs and develops AI/ML models, conducts experiments, and analyzes data for insights.	Data Science
Machine Learning Engineer	Bridges the gap between data science and production, ensuring models are scalable and performant in production.	Engineering
Data Engineer	Sets up data pipelines, manages data flow, and ensures clean, well-structured data for training and processing.	Engineering
Full Stack Developer	Integrates AI outputs with the frontend and backend, ensuring smooth user interaction and model integration.	Software Development
DevOps Engineer	Manages deployment, infrastructure, and **continuous integration/continuous deployment** (**CI/CD**) to ensure the AI models are smoothly deployed.	Operations
	NON-CRITICAL	
AI/ML Data Strategist/ Architect	Designs and implements data strategies to leverage AI/ML for business growth, aligning technology capabilities and business objectives.	Data Science
Data Analyst	Provides additional data insights, supports data quality initiatives, and helps with exploratory data analysis.	Data Science
UX/UI Designer	Optimizes the presentation of AI outputs, ensuring usability and enhancing user interaction with AI-driven features.	Design
UX/UI Researcher	Conducts user research to gather insights on user behavior, preferences, pain points, and expectations to help guide user-centric product design.	Design

Quality Assurance (QA) Engineer	Ensures that AI features are tested for functionality, reliability, and performance, reducing post-deployment issues.	Engineering
Cloud Architect	Helps scale AI infrastructure, optimizing cloud resource usage and handling large-scale deployment strategies.	Engineering
Business Analyst	Translates business needs into technical requirements, providing additional clarity and stakeholder alignment.	Analytics
Ethics and Compliance Officer	Assesses the model for ethical considerations, biases, and regulatory compliance.	Legal/Compliance

Table 6.1: Critical and non-critical roles

Now that you're familiar with what your dream team might look like, let's look at the next key component – your tech stack.

Investing in your tech stack

Understanding the tech stack and languages that will give you the most flexibility is the most important part of beginning your tech stack journey. As a PM, you're regularly thinking about the value of building something compared to the cost or effort required to build it. This is why you can't ignore your tech stack. It might be tempting to leave it to a CTO to decide on which tech stack to invest in, but it's not so easy to outsource if this tech stack is directly involved with the building of your product in any way. Empowering your stakeholders impacts the effort and cost that go into your AI product, so you'll want to make sure you're involved in the making of these key decisions.

You'll also want to build trust with your technical counterparts by developing an understanding of when your tech stack poses a limitation with your product. If you need to make a case for adding in more complexity for the sake and success of your product, or for the sake of customer acquisition, it should be you who brings it up. You'll work closely with your data science and data engineering teams to create the proper channels for delivering the relevant data to your models in a reliable way so that all the other stakeholders involved in building your product can trust the infrastructure in place. Here are some key factors that you should keep in mind as an AI PM:

- The biggest consideration in your tech stack is paying attention to properly setting up your infrastructure to handle the scale of your data as your business grows. Building the right tech stack to handle the cleansing, storing, securing, preparing, and monitoring of your data will be the foundation for supporting your AI program or product team. This goes hand in hand with building your data strategy.
- Beyond this, you'll seek to create a collaborative environment for the many roles in your AI program to participate in. Between the analysts, data scientists, ML engineers, broader engineering team, and any leaders and consultants you include in the formation of your AI team, you'll want to create channels for all these various members to communicate and collaborate.

- As is the case with many AI programs, you'll also be looking for ways to automate your processes and workflows. AI/ML technologies are consistently changing and improving, so you'll make a lot of traction when you invest in applying automation for your AI program itself. Use the brainpower of your talented staff to make higher-level decisions and use their complex problem-solving abilities. Anything that can be automated, refreshed, or standardized should be. This aligns with the precepts of agile and lean development. Refactoring and making elegant improvements in your code should be done proactively, so make sure to plan this into your sprints responsibly and use the advantages of your tech stack in your favor to set your team up for success.

While teams will work with R, Java, and C++ for ML, Python is by far the most widely used programming language for ML applications. A major reason for this is that it's a relatively simple language that's rich in libraries and tools to help with ML and data analysis. Python also has libraries like pandas and NumPy for tasks like data processing and manipulation, and libraries for data visualization like Matplotlib and Seaborn. There are also other popular libraries devoted to ML like TensorFlow, PyTorch, scikit-learn, and Keras to further expand on your ML models. Because Python is so widely used, an AI PM should be relatively familiar with it if they're going to work in the private sector.

Managing ML experimentation is a formidable undertaking in and of itself, and we've seen tools such as MLflow, Weights & Biases, and Neptune.ai that are used for managing versions and experiments. You can also use tools such as Cloudera Data Science Workbench, Seldon, Dataiku, DataRobot, Domino, TensorFlow Extended, Kubeflow, and SageMaker to support your data scientists with a workstation for building, experimenting with, deploying, and training ML models. Let's look at some tools that are most useful for complete model management, which includes quite a lot of aspects:

- When you're working with ML in production, you need to keep track of the versions of each model, along with that model's code, metadata, and configurations.

  Tools: **MLflow, Data Version Control** (**DVC**), and **Weights & Biases** (**W&B**).

- You're also packaging each trained model into some deployable container for when you're ready to deploy a particular model into production.

  Tools: Docker, TensorFlow Serving, and Seldon.

- Because you're often working with so many different models, you're storing all these versions in a centralized repository of some kind so you can easily retrieve different versions.

  Tools: **Model Registry** (**MLflow**) and Amazon SageMaker Model Registry.

- Depending on the models you're using and the demand for each, you'll need a more automated way of retraining the models you're saving on new data.

  Tools: Kubeflow, Airflow, and Tecton.

- We've discussed the concept of continuous monitoring in earlier chapters and certain tools also allow you to monitor ML models for performance and accuracy and enforce governance over policies and guidelines around the use of these models.

  Tools: Prometheus, Grafana, Evidently AI, and Neptune.ai.

- You can also keep logs of model performance, check system health, or deal with bugs.

  Tools: Prometheus, Grafana, and **ELK** (**Elasticsearch, Logstash, Kibana**).

Whether you're on **Amazon Web Services** (**AWS**), **Google Cloud Platform** (**GCP**), or Microsoft Azure, the cloud platform you're using will often dictate which products your organization ends up purchasing. These cloud platforms will often have their own tools for managing ML models themselves. Within this architecture, you'll also need containers to make your deployment scalable. Tools like Docker and Kubernetes can help with that.

Your specific needs will obviously depend on your organization's requirements, budget, and complexity. Our advice would be to only invest in more tools when you run into an issue with scaling.

The rise in accessible generative AI tools also plays a role in the tech stack since we're starting to see generative models play a role in optimizing the management of ML models themselves. Having enough data is a big strain on organizations working with ML models, so generative models can help with not only generating synthetic data but even editing existing data to make more diverse examples to bolster the amount of training data. Synthetic data can also be created for things like testing against edge cases or for certain biases.

We can also use generative AI tools like GPT-4 or BERT to handle exploratory data analysis by feeding them raw data and letting the models make visualizations, get insights, and offer generalized trends from the training data. Even the selection of the top elements in the dataset that contribute most to the models' performance, a process we refer to as feature engineering, can be handled by generative AI models.

Even the coding process itself can be optimized with the help of generative models. A company called Cognition recently announced a new generative AI tool called Devin, which can build, deploy, bug fix, and refactor code bases for various applications. It can also train and fine-tune ML models. Other products like Copilot and AlphaCode can also assist with this.

But don't go firing your staff for the sake of generative AI. These are tools and they're meant to be used as tools. You still need people to use them, to make incredible, complex things happen with them. Most AI/ML programs still fail even when they are properly staffed. Incorporating AI is not a silver bullet; it never was. The companies that treat it that way suffer for it.

Having the right process, talent, and tech stack to power your AI product is just the beginning. You're bringing all the necessary elements together to create AI-powered outputs. Those outputs will be integrated into your product somehow in the form of features that your customers will experience the benefit of. Next, it's important to understand how to orient the outputs of your AI/ML program and team in a way that's most evident to your customers.

Productizing AI-powered outputs – how AI product management is different

In this section, we will be exploring the difference between product management for traditional products and AI/ML products. At first glance, it may seem that AI/ML products aren't that different from traditional products. We're still creating a baseline of value, use, performance, and functionality and optimizing that baseline as best we can. This is true for every product as well as for the greater practice of product management, and this won't change just because our product works with AI.

The true differentiator when it comes to AI products is you're essentially **productizing a service**. Think about it for a moment. In order for AI to work, it has to learn from your (or your customers') data. Different models might work better on different kinds of data. Different datasets will require different hyperparameters from your models. This is the nature of AI/ML. In a way, this means that you could find yourself in a situation where the way you build and structure your product could even change fundamentally from one customer to another, especially at first when you have very few customers.

What this effectively means is you understand how all your customers are benefiting from the AI/ML *service* you're providing your customers and establish a standard procedure for recreating that process so that all your customers are able to experience the highest benefit. The general advice when productizing a service is first to find a niche or a specific cohort of customers to reach. We covered the UI/UX designer/researcher role earlier in this chapter; these folks will help you make the most of your customer cohorts and their preferences. If you can understand how AI is most helping your customers through your product and you're able to focus on what value means for your customers' success, you've begun the first steps of productizing.

The main focus will then be on how to best structure your product and internal processes so that you're not starting from scratch with every customer. This will be exploratory at first. Perhaps certain models are working better with some customers' data than others. Perhaps there are specific use cases that didn't arise until one of your customers asked for them. The basic idea of productizing is finding something that has value and brainstorming to understand where else you can leverage a previous project, output, or process for the betterment of all. Waste not, want not. If you can recycle something for the benefit of your collection of customers, it will add to the success and collective love your product receives.

Productizing services is a familiar concept in consulting where consultants are working project to project and often have to create new engagements with clients with an almost blank slate. But given that our own minds are pattern detection machines, eventually, we start to see similar requests come in or similar use cases pop up. Creating a repeatable process for certain types of customers, verticals, profiles, or use cases is the very essence of productizing. This is exactly what we're doing with AI/ML products. This is also a way to better anticipate how and why your customers are coming to you and your product specifically so that you know how to share successes and case studies that apply directly to them. The act of productizing will also immensely empower your marketing, sales, customer success, and broader GTM teams.

In many ways, the work of "productizing" ML happens within engineering, the folks responsible for building the application that will administer the models the data science and ML engineers are providing for engineers. You need engineers and developers whether you're building an AI product or not; these are the builders that turn a great idea, vision, or even roadmap into a tangible product you can experience. They're also responsible for making sure the product overall can scale, is robust enough to withstand a variety of user behavior, and ensures the product is compatible with various platforms and technologies.

As we have seen in this section, AI/ML product integration is about finding the balance between the outputs of your models and the use cases or groups those outputs are trying to serve. Many PMs come out of the experience of building with AI/ML finding they may need to customize their product for certain cohorts. So, next, let's address the need for AI customization, should it arise.

AI customization

Through the act of productizing, you'll likely start to create groupings or cohorts for your product. You might find that certain models work best for certain kinds of customers and build a structure around that. The act of grouping your offering into use cases and communicating differently or optimally to those cohorts builds on the wisdom gained from productizing. Taking that a step further, you'll then naturally start to verticalize.

Understanding the various considerations for business models, verticals, and special customer groupings will be an important part of how you go to market with your product. For instance, if you're supporting a **B2C or consumer product**, you'll want to invest more in information gathering by acquiring more of your end users' direct feedback for your ideation phase. Because you'll be creating a product that's going to be experienced by potentially millions of users, you will want a strong handle on the preferences and desires of those end users so that your product is most aligned with the *voice* of your customer. A product such as this could also benefit from experimenting and iterating fast so that you can get a product out and start gathering feedback on it before you're even getting your desired level of performance from the AI/ML models themselves.

If you're supporting **B2B or business-facing apps**, you might spend more time on the ideation and deployment of that product so that you can zero in on the specific set of customers you're hoping to win over. With business applications, you might be creating a very complex product that does a number of overlapping operations but for a very niche group of users. You might also have a wide variety of well-documented tools that offer similar solutions to your customers, so maintaining feature parity with the market could be more of a factor for this line of products.

The **size of your business** will also impact how you show up in your organization as an AI PM. For instance, if your organization is large enough, there might be a separate PM allocated to every version of the product cycle mentioned previously. There could also be various PMs that support various aspects of the functionality of your AI product. There could also be a PM that's devoted only to the AI features of the product with a higher-level PM or product director overseeing the entire suite. In a smaller organization, there might be only one AI PM overseeing the entirety of the product throughout all the phases of the lifecycle or ideating to continuously maintain that product.

Another differentiation between the sizes of companies lies in the **data** itself. Depending on how large your company is, the access or quality of the data you have can have varying degrees of difficulty. As a PM, making this data available and in a usable condition could prove more difficult and take up more of your time than it would at a start-up because of having to get access to siloed groupings of data that require administrative governance permissions. Even getting that data to fit together could be tough, depending on the conventions of each silo. Conversely, a start-up might have a hard time even amassing enough users and have issues with data volumes. Data availability and quality are notorious issues for companies big and small.

If you have a **highly verticalized product**, this could also impact how your work shows up as an AI PM. For instance, having a product that's specialized for healthcare, fintech, or education might necessitate a PM who has domain knowledge and expertise in these different verticals. Even if it's the same PM that's supporting this product for those verticals, you'll likely see it as three different versions of your product because the weights, thresholds, and performance metrics for these specific verticals will likely be drastically different. Depending on the type of product you're offering, you might even see a hyper-specialization to the point where the product is customized heavily even for every customer you have! This might be the case for highly tailored ML products that have to fit your clients' data.

Ultimately, the work of building products and crafting them for specific audiences is all about creating a product that sells. We need to be able to understand our audience—the buyers and users of our products—in order to build something that will truly serve them. In the next section, we will discuss the notion of selling and what that means from a product management perspective.

Selling AI – product management as a higher octave of sales

With AI product management, you're selling in a number of ways. There's the traditional sense of selling, which is this: creating a product that your market wants to buy. This is inherent in any traditional PM role. Your understanding of your market, your product, and your sales force is one where you're confident in the solution you're bringing to market, your solution's performance, and your sales force's ability to articulate the value of the solution you're building. Then, there's the opportunity and challenge AI presents: the ability to sell the AI functionality itself.

Earlier in this book, we mentioned the difficulty AI/ML projects have in being deployed into production; this happens for a number of reasons, but among the top reasons is the inability to sell the value of AI to the broader organization. This will be incredibly relevant to any PM who's looking to work with AI products because you will have to develop this understanding and ability to sell your product to the outside market as well as to your internal stakeholders and leaders. This is the case whether you work at a large organization or a small start-up.

Part of the work of selling AI within your organization is going to directly correlate with how you empower your teams. If you want to be successful in evangelizing your AI product, we have the following advice:

- Make sure there's alignment on the data strategy and data architecture side. Remember: AI products are data products, and if you don't have sound practices on the data end, you'll have to reverse engineer your data pipelines, and that's not fun. Part of this is also making sure there is a sound data strategy that will grow as your business and product scale to support more users and customers.
- Make sure that you are not focusing so much on the technology and tech stack that you're forgetting about tying the use of that technology to impact. Remind yourself why you want to use AI for your product. What value is it providing your product and, by extension, your organization and customers? No sale is successful without perceived value. People who work in sales understand they need to be successful in communicating the value of their product before you'll open your wallet. You'll have to have the same mentality if you're working in product management. This isn't just relevant for your customers but for your leadership team as well.
- Make sure you're embracing experimentation. As we've said countless times, AI/ML is a highly iterative process. You can't load up your practitioners and the models themselves with expectations and results without first allowing them the space to experiment. Beautiful, miraculous things can happen with AI/ML, but first, they need to be coaxed through iteration. Flexibility and curiosity will also be important when you're considering various ML models that will support your product as it scales.

As with other product roles, the role of an AI PM is vast and requires knowledge and competencies in several areas. Having an understanding of how an AI product is created and brainstormed allows you to meet your market's needs with leading-edge technology. Knowing how to best support the data end of such a product means you are setting up the most integral members of your AI team for success. Allowing your hands-on AI practitioners the space and ownership to experiment and report their findings to you and your leadership creates the proper feedback loop between tech and leadership to see your product flourish. Finally, creating the proper deployment tech stack and process ensures your end users are benefiting from your product in the way you originally intended and gives you the perfect board with which to dive into the day-to-day process of iterating on your product until its sunset days. Let's illustrate these points with a quick case study.

Case study

In this section, we're introducing a case study example that we'll be discussing throughout *Part 2* as a working example of how the concepts and principles are applied. In this chapter, we will be getting an introduction to a sample company and product, the way they organize their product organization and cross-functional teams, their tech stack, their productized outputs, how they've customized their verticals, and a brief overview of their GTM strategy.

Waterbear Inc. is building a flagship product, Akeira, to help tackle the mental health epidemic we've seen been exacerbated by the Covid-19 pandemic, specifically for women over the age of 18. What began as an idea to help people better understand themselves has now turned into a viable product people are using to get a better read on how their sentiments and moods are changing over time. The product uses NLP to scan journal entries that users type into the application, establish goals and focus areas the user wants to work on, and recommend activities and questions to better help them reach their personal goals.

Despite funding hurdles and difficulties early on to acquire users, they've now reached a milestone. They've acquired series A funding and have now acquired the resources needed to pursue new customers and build out their internal teams. Their mission of helping individuals better understand themselves and rise above their psychological blocks is well on its way. Waterbear's approach is holistic; they want to build an organization where job functions are seamlessly integrated to ensure not only an efficient product lifecycle but also a healthy pace for the company's scalability and employee well-being.

Based on concepts we've discussed in this chapter, we will be offering a breakdown of how the product navigated through each phase of the AI product development cycle, how various job functions work together to run this product, what their tech stack looks like, which AI outputs make the most impact on the product, how they've verticalized this product so far, and what their current GTM strategy is.

AI product development cycle

First, let's revisit the stages of the product development lifecycle covered earlier in this chapter and see what this looked like for Akeira:

- **Ideation**: The Akeira product team first began by analyzing how users might find value in a mental health tool to offer real-time emotional support, analysis, and reassurance. AI features like sentiment analysis, **natural language understanding** (**NLU**) for therapy sessions, and mood tracking through conversational data were all considered. Feasibility studies were conducted to make sure the data requirements aligned with potential solutions and models that would empower the kind of complex language understanding required for mental health contexts like detecting distress or tone.
- **Data management**: Mental health-related conversation datasets and anonymized therapy transcripts, mood-tracking journals, and relevant social media interactions were acquired to support training AI/ML models being used. Emotional tones and sentiments were categorized and labeled in preparation for feeding into the AI/ML models that would be tried and tested in the third phase.
- **R&D**: Various language models were used but the team finally settled on Google's BERT model. The collected and labeled mental health data from the prior phase was used to fine-tune BERT for specific tasks like detecting mental health cues in conversation. BERT was then tested for performance thresholds in understanding context, intent, and emotional sentiment before being integrated into a working prototype or MVP, which became referred to as Akeira. This prototype was then shared with a small group of beta testers.

- **Deployment**: Akeira was incorporated into a mental health app of the same name and monitoring tools were set up to track how well it was able to identify emotional cues in real-time interactions and respond. Now that beta testers were actively using the app, teams were able to set up maintenance and monitoring to ensure the beta testers were using the app correctly and were seeing benefits from it. Feedback mechanisms were put in place to capture their concerns and issues. These feedback mechanisms became a prime feature of the app to ensure Akeira could regularly gather feedback about the accuracy and sensitivity of Akeira's reactions so that future iterations were adjusted accordingly.

Team breakdown

Let's take a look at what the team working on Akeira might look like:

- **Product**: One of the first hiring decisions the company made was to find a worthy hire for their coveted founding lead PM role. They were looking for someone who had experience building AI tools focused on NLP, but they were also hiring for passion. Once they found someone who had a proven record for building products that help people, with a personal interest in psychology and well-being, they were able to trust that person to define the strategy, vision, and roadmap for the company's flagship product.

  Because this person has built AI products in the past, they're able to work well with the subject matter experts, therapists, data scientists, ML engineers, testers, customer success folks, and designers to bring their product to the next level. They're able to combine these internal workings with research, competitive analysis, user feedback, and external data sources to understand where there might be opportunities for the product moving forward. They understand that their objective is for their product to not only remain relevant but also to innovate and align with the needs of the market they're serving.

- **Data science and ML**: The data science and ML team serving the product has their work cut out for them. These folks are responsible for building and training the models that power Akeira.

  The data scientists are routinely gathering insights from their user data to better understand what kind of user behavior leads to the optimal use of the product. Their work focuses primarily on A/B testing various models and using statistical, probabilistic, and predictive approaches to NLP to continuously improve their AI models, and they regularly share their findings with their product and ML counterparts. They also spend a lot of time bolstering various **large foundational models** (**LFMs**) with specialized, technical data that comes from the therapists they work with to ensure the viability of the product. Most of their time is spent on exploring the data, feature engineering for the models, choosing the best models for deployment, and using their knowledge and expertise to interpret the results of their models. They're primarily coding in Python and they hand off these models as outputs to the ML engineers they work with.

In contrast, the ML team is more focused on building, deploying, and maintaining the AI systems that run Waterbear's product in the company's production environment as a whole. While they may get in the weeds for hyperparameter tuning or making adjustments to the models, this isn't the focus of their role. The main focus is on optimizing the ML pipelines that allow the product to be scalable, something that will be increasingly important as the company grows its user base. Once the ML engineers get the models developed by the data scientists, they're responsible for productionizing them. This means they're responsible for making sure they run smoothly in production and they work closely with software engineers to ensure the models are integrated into the rest of the product experience for the end user. At Waterbear, these folks work very closely with data science and engineering counterparts to ensure product expectations are met by both.

- **Frontend/backend/full stack engineering**: At Waterbear, they're responsible for dealing with breakdowns in the code base and refactoring. Any feedback or bugs that come back from QA and testing at the functional level would fall in the domain of engineering. They're also trying to make sure the product is able to scale functionally, and that the company has invested in enough computational resources to handle the demand brought on by the ML models and data load. They're also responsible for making sure the product's overall performance and speed are optimized, that it's safe from data and cybersecurity risks, and that it's integrated into the cloud services and platforms the company has invested in.
- **QA/testing:** Because Waterbear's product is a B2C product, the QA and testing team spend quite a lot of time creatively coming up with ways users might interact with Akeira that fall outside of the intended use for the app. These folks not only work on testing strategies but also ensure the product is unit-tested. This means that all the functions and components of the app work as intended. They also test to make sure all the integrated parts of the product are working as intended too, that the data inputted into the app finds the right path to the dashboard someone will see as a user, and so forth.

  Finally, another crucial part of the product that needs to align is that the user criteria that were part of building the product out in the first place are still being met. Because all product decisions have a user-specific reason for making it to the development sprints, all work that is output by those teams needs to meet the criteria for being prioritized in the first place. In this product organization, the testers ensure that happens as a catch-all in case the engineers have missed something, so they work heavily with product and engineering teams to make sure goals and objectives are met at the granular level.

- **Customer success**: This role looks a bit different at Waterbear because, in this case, the customer success team isn't working directly with customers but they're staying aware of trends that come up in the overall customer base. What this looks like here is keeping an eye out for unintended use cases, trends in frustration across the customer base, and opportunities for new features as well. Data was obtained through various focus groups and product surveys for beta testers.

At Waterbear, this role is primarily focused on regularly delivering insights to their counterparts on the product team and serving as the company's "Voice of the Customer." Customer success specialists are compiling trends from the insights that come from direct customer feedback and are discussed regularly with the product, UX, and design teams at Waterbear.

- **UX and design**: Crafting an intuitive user-friendly customer experience is the main objective of these teams. At the UI level, a lot of problems can arise as users are navigating the app that product, engineering, ML, and data teams might not even be aware of. This is where UX and design come in. For Waterbear, these roles work primarily on minimizing frustrations and maximizing delights for the end users as they experience the product. While data scientists are A/B testing the functionality of models, these folks are A/B testing various stylistic choices that are more visual and experiential in nature. Their main goal is to design and brainstorm a product experience that's both visually appealing and compelling for the end user.

At Waterbear, none of these roles work in a vacuum and the company works on integrating all of these main areas through the use of regular product strategy meetings. Reinforcing the product vision and strategy is a big part of the product team's responsibility at Waterbear and what that means is having all teams represented in some way to ensure each focus area has a voice. Each team has a lead and that lead is a crucial member of the product strategy meeting.

Throughout *Part 2*, we will see all these roles start to change and evolve. As we add more complexity and definition to the product and operations at the company based on the concepts we discuss in subsequent chapters, this case study example will grow with us.

Tech stack

Waterbear's tech stack that supports its Akeira product is based in the AWS ecosystem. Because they were initially accepted into an Amazon start-up accelerator pilot program, they were able to scale and build their products with Amazon services at a low price as they acquired more customers. The following is a breakdown of what their tech stack looks like:

- **Product:**
  - Jira for managing engineering tasks and analyzing the sprint cycles as well as tracking bugs
  - Confluence for documentation and knowledge sharing
  - Slack for communication and collaboration
  - Mixpanel for analyzing user behavior in-app
  - Hotjar for heatmaps and bolstering where Mixpanel falls short
  - Figma for making mockups and wireframes
  - DataDog for infrastructure/application performance testing
  - ProductPlan for managing a roadmap

- **Data and ML:**
    - Amazon S3 buckets for storing vast swaths of data
    - Amazon Athena for querying and analyzing the data they have
    - Apache Spark for the pre-processing and feature engineering of their models
    - Jupyter Notebook for their own R&D of ML models
    - GitHub for managing versions of models
    - Amazon SageMaker to fully manage their ML pipelines (building, training, and deployment)
    - SageMaker Ground Truth as a labeling service for all the components that relate to the user dashboard to capture things like entity recognition, text classification, and sentiment analysis
    - SageMaker Clarify as a de-biasing tool for their NLP models
    - Amazon **EKS** (**Elastic Kubernetes Service**) for containerizing all NLP applications relevant to Akeira
- **Engineering:**
    - AWS Lambda functions to handle the pre-processing and post-processing of new data from users to feed into the models in production quickly
    - Amazon API Gateway for handling the APIs that connect the NLP models that power the product
    - Amazon CloudWatch to monitor performance logs
    - AWS X-Ray to debug issues with NLP services
    - AWS **KMS** (**Key Management Service**) for data protection and encryption of the data
    - AWS **IAM** (**Identity and Access Management**) for controlling access to specific users
    - GitLab for code management
    - Jenkins for CI/CD
- **Other teams:**
    - Intercom for user feedback, support tickets, and in-app surveys
    - InVision for designing one-pagers, product marketing material, and white papers
    - Google Forms for user surveys
    - Lookback for conducting more formal surveys en masse

AI outputs

With regard to Waterbear's Akeira product, user journals are recorded and integrated into a dashboard the user is able to reference as they're going about their day. The following is an example of what that might look like in the app:

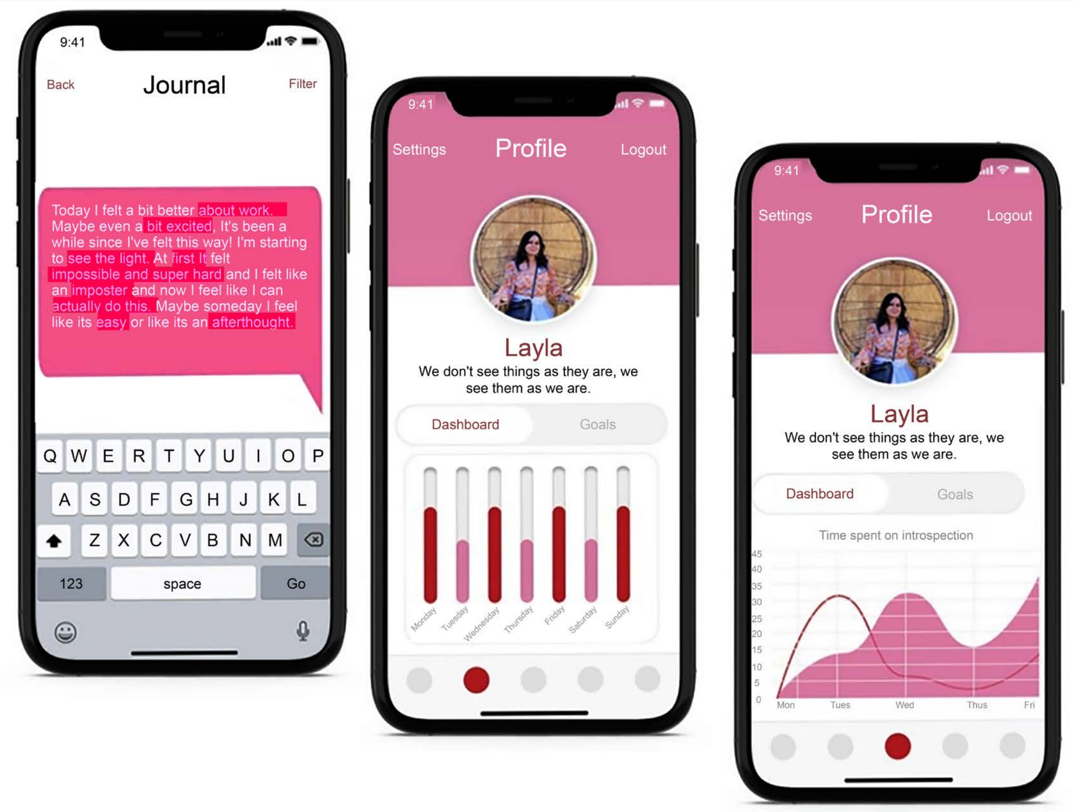

Figure 6.2: Akeira dashboard with journal and personal data

Because the app is a consumer-facing mobile app, the text that would be shared by an end user would be analyzed and keywords would be selected. The app would first aim to understand the end user in the form of a chat, so Akeira's first objective of making sure the person they're talking to is seen, understood, and validated is achieved. NLP models would analyze not only the text explicitly shared but would also be connected with social media and other text inputs if the user wants it to establish a holistic understanding of how they're doing. These would be the inputs.

The outputs that a consumer would see would be shown between a toggle between the dashboard and the goals modules. The dashboard module would be able to show various metrics you might use to gauge your overall health and well-being. An overall Akeira score would be shared as a visual indicator of your overall emotional state over a specific period.

Other KPIs and visualizations would also be included in the dashboard. Here, you can think of things like a mood tracker, engagement levels on the app and how they correlate to your mood, self-care metrics with progress bars, visualizations relating to physical activity, a gratitude score to reflect areas where the end user can build appreciation for the areas of their lives that are going well, and a social tracker that would allow users a bird's-eye view of their social health:

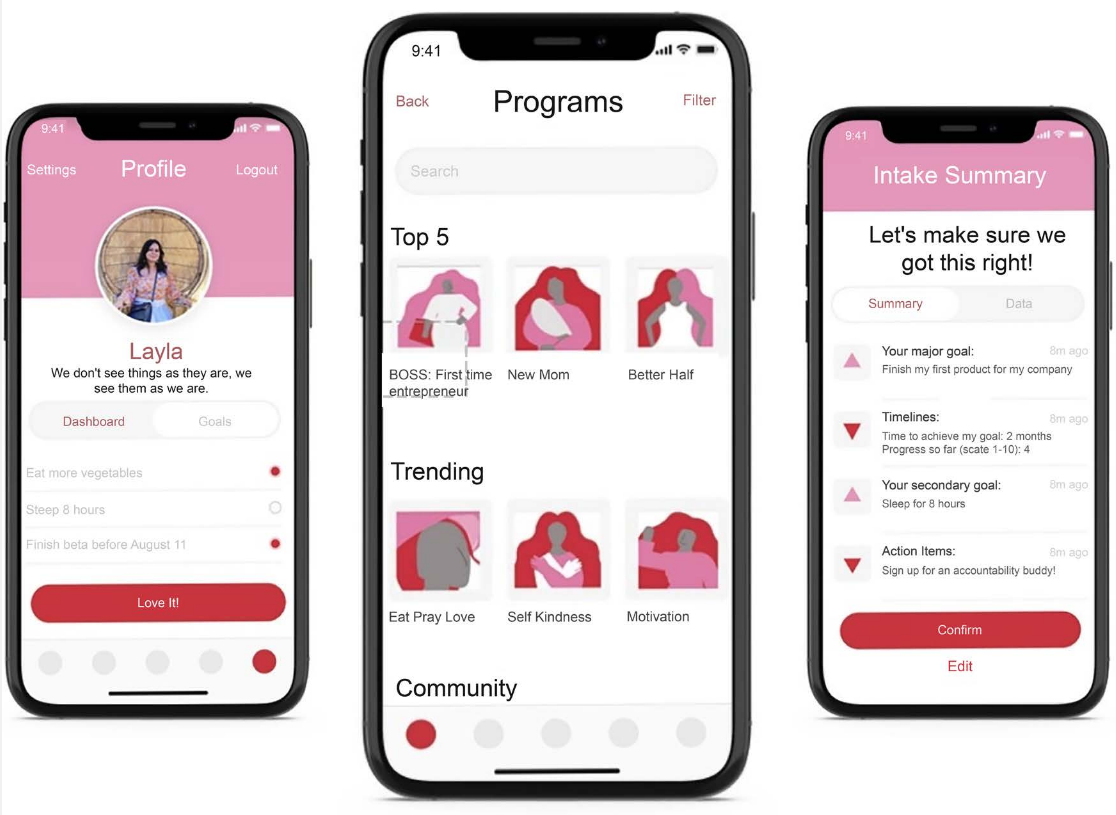

Figure 6.3: Akeira dashboard showing the available programs

The second toggle would be their established goals module that specifically covers goals that were established through the NLP models and confirmed by the user; these would be more specific in nature and tailored to the end user's specific use of Akeira. Goals would have established success criteria so that once progress is made toward a goal, it can be recorded and celebrated as well.

Finally, all these various metrics and KPIs would result in highly personalized insights for each end user for things like feedback display, prompts that encourage self-reflection, mindfulness, or even actions for physical and mental well-being. If the goal of an AI product is to deliver value through data, the app experience needs to have a way of expressing that value and reflecting it back to the end user. This is the way Akeira and Waterbear communicate with their customer base.

GTM strategy and verticalization

Waterbear initially opened up to a small group of beta testers early on to better understand interest and viability in a product like Akeira after completing the prototype. The target market was initially adult women over the age of 18 and focused on authentic and organic marketing campaigns that met their users where they were: endlessly scrolling online. Because both of Waterbear's founders were women, they were able to build an early network of potential investors and advisors through a start-up accelerator with Women in AI.

The market for a product like Akeira was also a blue ocean market. While there were direct competitors like Woebot, Wysa, Moodnotes, and Pacifica out there, most of them were simple mood trackers or chatbots that didn't quite capture the complexity that Akeira was able to offer. Other indirect competitors like Talkspace, BetterHelp, and Cerebral were only able to get you connected with mental health professionals in the space. Because Akeira was a combination of both direct and indirect competitors, it was able to amass a loyal following of dedicated users using a product-led approach.

Using conservative estimates of the **total available market** (**TAM**) of 0.6 billion, based on the state of the mental health-focused app market, Waterbear was able to discover a shared **serviceable addressable market** (**SAM**) of mental health apps focused on women to be just over $350 million. Because they prioritized slow, organic growth with their one flagship product before focusing on scaling, their initial **serviceable obtainable market** (**SOM**) was just over $5 million.

Finally, although Akeira was initially created to service direct consumers specifically, Waterbear's GTM team discovered that there was a significant population of users that could come from businesses that were looking for a conversational AI solution they could deliver to their employees to help boost their creativity, happiness, and well-being.

Future explorations of Akeira would branch out into other verticals like various age groups, genders, and use cases beyond just well-being. Look out for our Waterbear case study through the rest of the chapters in Part 2 as we add more complexity to our example.

Summary

The work of a PM is never done. There are always more voices, perspectives, and considerations to take in. Coordinating all the stakeholders, technology, leadership, market analysis, customer feedback, and passion for a product isn't an easy task. In this chapter, we covered the stages of the AI product development lifecycle and the various roles that can make up your AI product dream team. We also covered the tech stack that can help that team build a product, and various focus areas to help that product stand out and resonate with the cohorts of groups that will be buying and using your product. Hopefully, this chapter has helped you understand what the most important factors are when you set out to build an AI-native product.

As long as you're hiring the right people for the roles you have open in your AI program, doing your due diligence to uncover the right strategy for tech stack adoption, structuring your product in a way that benefits your customers according to their verticals, and working with your leadership and greater GTM team to build a product that meets a need in your market, you'll be set up for success. You're likely to do better than you think.

In the next chapter, we will go deeper into the semantics of building a product versus productizing an AI/ML service. We will go over the differences between AI-native and traditional software products, productizing under various business models, performance evaluation, and customer feedback loops.

References

- *The what, why, and wow of A/B testing in machine learning*: `https://mlops.community/the-what-why-and-how-of-a-b-testing-in-ml`
- *TFX*: `https://www.tensorflow.org/tfx`
- *Seldon*: `https://www.seldon.io/`
- *Dataiku*: `https://www.dataiku.com/`
- *DataRobot*: `https://www.datarobot.com/`
- *Domino*: `https://domino.ai/`
- *Cloudera Data Science Workbench*: `https://www.cloudera.com/products/data-science-and-engineering/data-science-workbench.html`
- *Weights & Biases*: `https://wandb.ai/site`
- *MLflow*: `https://mlflow.org/`

7

Productizing the ML Service

In *Chapter 6*, we briefly touched upon the notion of productizing and what that means for AI outputs in the *Productizing AI-powered outputs – how AI product management is different* section. We will be expanding on that concept in this chapter by exploring the trials and tribulations that may come up when building an AI product. Rather than thinking of AI products as traditional software products, it helps to think of them as a service that you're learning to productize. **Productizing** is the process of taking a service, process, skill, or idea and finding a way to present it to the greater market.

First, we will look at some of the basic principles of productizing and the differences and similarities between AI products and traditional software products. While we've covered some of these elements in prior chapters, we will be expanding on these ideas further in this chapter. Next, we will explore how productizing might change based on the business model itself. We'll take a look at the role **AI for IT operations** (**AIOps**) and **machine learning operations** (**MLOps**) play in productizing when it comes to managing your various products and comparing their performance. Finally, we'll end with a look at how our productizing loop can be continually refreshed through feedback.

By the end of this chapter, you will have an understanding of the following topics:

- Productizing principles
- Understanding the differences between AI and traditional software products
- B2B versus B2C – productizing business models
- Consistency and AIOps/MLOps – reliance and trust
- Performance evaluation – testing, retraining, and hyperparameter tuning
- Feedback loop – relationship building

Basics of productizing

Productizing is a concept that comes up a lot in consulting, where service offerings are solidified into well-defined entities that can be effectively customized and resold. What this refers to is the ability to create a consistent workflow that you can rely on to deliver consistent results in the way traditional products demand.

Before we go further in depth into product management principles and aligning them to the idiosyncrasies of AI/ML services, let's address a few of the basic ideas when it comes to what productizing means in the first place. The following flowchart demonstrates the process of "productizing" through four major areas: defining the scope of the product you're ending up with, aligning branding and packaging to support that scope, training your sales and delivery teams to appropriately communicate the scope of your product, and establishing a feedback loop to inform the scope of your work (or to establish new productized services).

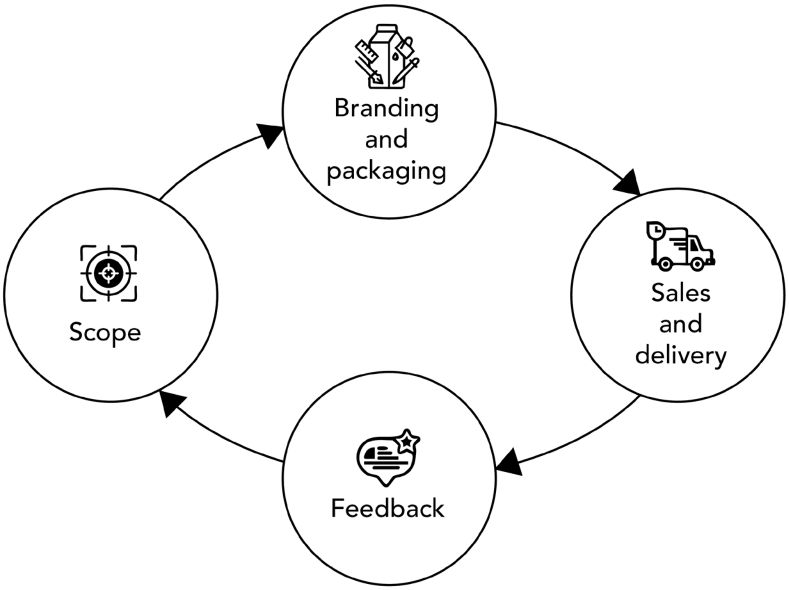

Figure 7.1: Productizing basics

Because the AI/ML models and outputs are what are effectively being productized, this means that standards will need to be set for the following, as shown in the preceding flowchart:

- **Scope**: Defining a scope establishes boundaries of what your product will offer to your customers and users. This is where you define the features, functionality, deliverables, and outputs your end users can expect. The scope also guides internal teams to better understand priorities and expectations. It also helps create a roadmap for future iterations and variations. Some questions you can ask at this stage are:
    - What are the features, deliverables, and outputs you are offering as a product?
    - What can your customers consistently expect from your product?
    - Will you have different models for different customers?
    - Will each customer's data stay separate to train on separate models or will you train your models on all your customers' data at once?

- **Branding and packaging:** Whatever scope you set will inform your branding and packaging choices. If you are grouping models and deliverables into cohorts, you'll effectively need to brand those as separate products, each with its own marketing collateral (whitepapers, one-pagers, blog posts, etc.), contracts and **service-level agreements** (**SLAs**), documentation, brand identity, and messaging. Some questions you can ask are:
    - What is the core value proposition of this product?
    - How will the brand voice align with this product?
    - How is this product best visually represented?
    - Do branding and marketing efforts align with the target market for this product?
    - Do branding and marketing efforts need to be tailored across a variety of target customer profiles?
- **Sales and delivery**: Your sales force will need to be trained on your various products, however they shake out. They'll need to understand how to talk about each product and the competencies and value propositions of each. You will have to agree on a method of delivery here as well. Some questions you can ask are:
    - Is there a platform or portal where each product will be administered and delivered?
    - Is it an application on a mobile phone?
    - What level of automation will you be embedding into your delivery architecture?
    - How will the product be priced? Will there be tiers?
    - What are the key pain points addressed in a sales pitch or deck?
    - How is value going to be communicated during the sales cycle?
    - What will the onboarding process look like?
- **Feedback**: You'll need some way of measuring how your product is resonating with your audience and cohort group. You'll want to do this to make sure you have happy customers whose needs are being met by your product but you'll also want the feedback to be incorporated into the future *scope* of your product as well:
    - How will customer feedback be gathered, used, and analyzed?
    - How will negative feedback be addressed and prioritized for future iterations and product updates?
    - What communication plan exists for sharing improvements and updates with customers?
    - Which customer success metrics will be tracked?

Now that we've gone over some of the basics of productizing, let's explore how these aspects might change based on context.

AI versus traditional software product management

Understanding the differences between traditional software and AI products isn't so much about comparing two disparate groups of products. Although this is helpful, this is a macro trend, and the biggest reason for understanding the two is to anticipate how all products will evolve with AI. From the perspective of a PM, there are a number of differences between traditional software products and AI/ML products. In this section, we'll first go over how they're different, and then we'll note the similarities between the two to give us a well-rounded sense of what to expect when you're product-managing an AI/ML product. This will help us establish a baseline as well as a deviation from traditional PM work.

How are the products different?

While AI products are built on a foundation of traditional software development for the most part, they do have a number of key differences you should be aware of as a PM. AI is able to evolve a traditional software product, and you'll hear this referred to as applied AI in product circles. What this refers to is applications of AI outside of a research setting or lab that are used in the building of tech products. Essentially, the concept of putting AI/ML to use, testing and optimizing the models for accuracy and precision, and evolving them over time through feedback loops is what constitutes applied AI.

The following sections will cover the biggest differences between traditional software products and AI products, which include scalability, profit margins, and uncertainty.

Scalability

One of the major differences between applied AI products and traditional software products is in the area of scalability. Because AI is so specific and sensitive to the quality and peculiarities of the training data, you're likely going to have issues with scaling this kind of product because you're likely to encounter so many edge cases that you have to go back to the drawing board or start to create cohorts within your user base. This has led to what AI Forum refers to as a collective AI fatigue from the people who are working directly on developing AI, as well as with leadership and the market at large itself, that all stems from the issue of scalability when applied for commercial uses.

This issue lies not just in the challenge AI poses with accuracy; indeed, you need a lot of data for advanced ML and **deep learning** (**DL**) models to work well and give you good performance, but you also need those models to be robust enough to work with a generalized population of users, situations, contexts, and locations. In traditional software products, you might not need this level of granularity to see success, but with applied AI products, you might think you need a highly personalized model for every user or use case.

This sense might not be that far off from an optimal reality. What this also means is that AI is particularly sensitive to the idea of edge cases, and the threshold for what even constitutes an edge case might be lower for AI products compared to their traditional software counterparts. While edge cases impact all software products, traditional software companies do have an inherent advantage in that they're iterative, so once you build and ship a product, you're able to sell it to virtually all your customers without having to do a lot of customization.

We will go over the differences in business models later on in this chapter, but this issue of scalability is doubly affected by your choice of business model. For instance, if you're a B2B company, your AI product might behave wildly differently when compared to other customers because their training data differences might be too great. If we contrast this with a B2C company, you might be training your models on really representative and diverse data, but then the way your individual consumers interact with your product might vary wildly.

Whether an issue with the training data or the way end users work with your product, the issue is the same: you're having to account for so many perspectives and demands on your product from outside influences that it makes the scaling of the models used almost impossible while keeping with one consistent product build. Even the issue of agreeing on an acceptable level of product performance will likely be time, cost, and energy inefficient, let alone actually acquiring those levels of accuracy that you'll need to end up with a product that's consistent enough to sell as an a minimum viable product.

Getting to a level of trust where your earliest customers will see the value in your product enough to consistently use it will require a lot of initial work, and you don't have guarantees that this intensive customer acquisition will necessarily dissipate as it might with traditional software products. With applied AI products, the process of acquiring and keeping your customers may stay at a consistently grueling pace, which further contributes to this issue of scalability up to a certain point. Strides are being made, however, to build and discover ways to improve models that are less reliant on massive hoards of data.

According to AI Forum "Data can be expensive to collect, requiring annotation and labelling (by clinicians in healthcare), cleaning, and preparation which can contribute 50% of AI training costs. Consent from data owners (e.g. patients) is needed to use their private data, and additional incentives are sometimes required for data custodians to share the data. Data privacy and security laws can introduce barriers to sharing, storing, accessing and handling the data."

This reinforces the idea that the quantity of data itself is secondary to the commercial success of scalable AI products compared to the quality and diversity of the training data. It doesn't just need to be standardized in a way that's uniform across your data sample; it also needs to have enough representative data points that encompass the diversity of users as well—the idea being that the more diverse your data is, the more it will be able to anticipate the needs and uses of the general users that will experience it once it is deployed to the greater public.

Access to clean data is easier said than done. Most real-world datasets are riddled with data hygiene issues. Tableau (`https://www.tableau.com/learn/articles/what-is-data-cleaning`) offers a great summary of how to solve this: "Data cleaning is the process of fixing or removing incorrect, corrupted, incorrectly formatted, duplicate, or incomplete data within a dataset. When combining multiple data sources, there are many opportunities for data to be duplicated or mislabeled. If data is incorrect, outcomes and algorithms are unreliable, even though they may look correct. There is no one absolute way to prescribe the exact steps in the data cleaning process because the processes will vary from dataset to dataset."

AI/ML products are particularly sensitive to the quality of data, and this is further exacerbated by the issue of deliberate tampering or data poisoning in which end users attack AI/ML systems to intentionally manipulate the training data that powers them.

Profit margins

We've discussed the costs associated with building out an AI organization, which you'll have to do if you're building applied AI products. These costs are some of the biggest differentiators between traditional software products and applied AI products. Because the development, data, and maintenance costs can be so high, they will impact your margins. To paint a picture of what can go wrong, the best example we were able to find came from Harvard Business Review (`https://hbr.org/2022/03/how-to-scale-ai-in-your-organization`), which mentioned "...one financial company lost $20,000 in 10 minutes because one of its machine learning models began to misbehave. With no visibility into the root issue — and no way to even identify which of its models was malfunctioning — the company was left with no choice but to pull the plug. All models were rolled back to much earlier iterations, which severely degraded performance and erased weeks of effort." Imagine losing that much money in such a short time and still not being able to know where the problem was!

With that said, there are aspects of applied AI that are more profitable than others, depending on the timescale:

- According to Forbes (`https://www.forbes.com/sites/forbesfinancecouncil/2020/09/01/past-the-ai-hype-when-is-ai-actually-profitable/`), the area where we see the fastest short-term return on applied AI is when it supports customer behavior, which makes sense because it directly impacts sales and is one of the highest revenue drivers. This would be use cases of applied AI, such as product recommendations, scaled pricing algorithms, or advertising personalization/optimization, that bring in more traffic to your website or help customers choose more products based on their specific spending and purchasing habits, cost consciousness, or tastes and preferences.
- The more medium-term payoff of applied AI is when AI is used to improve a product or make a user's experience better in some way. This could look like product automation that boosts users' productivity. For most of this book, we've been talking about these applications of AI because they directly impact an individual product and its performance. As a PM, this is likely the primary area you are focused on when it comes to applied AI and where the best place for it to fit is as far as your product's performance and standing with customers and the wider market are concerned. Because this essentially poses a medium-term payoff, and because your AI costs will be front-loaded for the most part, you'll have to manage expectations with your engineering, leadership, and customer-facing teams from both a revenue and performance perspective.

- The most significant long-term payoff of applied AI is when it affects the reputation of the company. This means AI is applied in a way that boosts a company's trust reputationally, particularly when compared to its competitors and peers in the same space. But acquiring this level of AI supremacy is time intensive, and it sheds light on the double-edged sword of AI. Eventually, your competitors will catch on and will pay to play in your arena as well. This is the nature of competition, and as PMs, we understand that our products are scrutinized in the market, and rightly so. That being said, these areas of profitability aren't exactly siloed. The short-term advantages of AI that we just discussed can also bleed into this area if your product specializes in marketing, recommendation systems, or advertising, for instance.

  The long-term advantages of AI that we discussed can also apply to integrating applied AI into your product if the premium AI offers your product is so great that it impacts your reputational standing in your chosen market or vertical. As a PM, technologist, or entrepreneur, you're going to have to grapple with the costs AI poses as well as its advantages and create a plan for how you want to start leveraging applied AI. As you're getting started, one of the best ways to do this is to first understand the benefits AI can offer your product and company at large and then draft your own version of this plan and pitch it to your leadership. Communicating the effectiveness of AI and your understanding of how it will impact profitability will help you gain credibility in your own organization. It will also offer an avenue for your stakeholders to be able to understand the challenges and opportunities and ultimately contribute to the plan.

Once you have this plan, you can start keeping track of the relative profitability of the various areas AI is applied to and have champions/supporters within your organization that help keep the visibility of AI. If you can all agree on a specific appetite for AI spending and margin threshold tolerance, it will be easier to navigate the AI waters as you continue on your applied AI journey because it will force you to invest in AI in cases where there's data, economic returns, and excitement to keep it going. It will also keep your finance team happy.

Uncertainty

The last big difference is conceptual. AI introduces a whole lot of uncertainty to the work of product management. For a PM, this may be the most important section of the book thus far. A PM applying the known patterns of product management from e-commerce, banking, travel, and social media will only take them so far when building an AI product. PMs need to adapt their thinking and their approach in light of the uncertainties AI can introduce into their product.

With traditional software products, the deterministic qualities of a product are hardcoded. Algorithms still exist, but they're not improving or learning over time; they are static. But with AI products, these qualities are more fluid. There's a level of expectation setting and performance that has to be agreed upon, expressly or intuitively, by the builders and users. If that level of performance doesn't happen because the models aren't trained enough, they have to be trained more. If performance doesn't come no matter how much training you throw at it, you might not have a product to sell at the end of the day! This level of uncertainty is cushioned in traditional software because your performance goals aren't quite the moving target they are when you're building an AI product.

Some factors you may have to keep in mind are:

- Where ML is concerned, you're bound to have some level of error because, as we know, no model is going to be perfect, particularly if you're using generative AI models, which might introduce factually incorrect information. As a rule of thumb, more data will often get you better performance. Collectively with your leadership, customers, and engineering team, you will come to understand where the threshold of accuracy needs to lie. It will be different for every product and every use case, so this is something you'll need to discuss and come to an agreement on.
- The output from one model based on a certain training dataset might be great one day, but if you diversify the training set and retrain, it's probable you won't get the results you're expecting. It's very hard to exactly recreate and test these products, and keeping an experimental attitude when it comes to managing how you are testing, deploying, and maintaining versions will be essential. It's also hard to know whether your models will significantly improve because of the type of model you've chosen, the training data selected, or the features you've selected.
- You also don't know how much time, how many resources, and how much data you need, which makes it really difficult to actually plan your roadmap and keep a sense of time in the way you might be used to if you've worked with traditional software products. If you've been a PM before, that last sentence might raise some eyebrows. Yes, indeed: AI products are even more difficult to forecast! Finding the right balance of factors might take days, weeks, or months. With AI, it might not be until you're well within the weeds of building that you start to grasp the complexity of your scope. This poses a great challenge to your leadership team, which might be interested in more concrete answers about the length of time you need. AI introduces so many factors of uncertainty to keep track of, which is why it's so important to find a leadership team that gets this and supports you because the journey is riddled with doubt.

All in all, AI is here to stay, and it's going to be a major theme as we head into the 2030s. The promise and opportunity of AI continue to grow as we see companies applying AI in new and innovative ways, and our guess is as the decade continues, we will see more inspired ways for companies to benefit from AI adoption. Companies investing in and building AI-native products today will be setting themselves up for success for decades to come if they survive. The key is surviving. According to Google's CEO Sundar Pichai, "the risk of underinvesting [in AI] is dramatically greater than the risk of overinvesting" (`https://www.cnn.com/2024/08/02/tech/wall-street-asks-big-tech-will-ai-ever-make-money/index.html`).

But this is a tricky concept that even the biggest tech companies struggle with. We can look at a more well-known example: Tesla. Tesla has arguably one of the most potentially lucrative applications of AI on the market: **Full Self-Driving** (**FSD**) functionality, an upsell feature that has direct revenue implications that's been on the market since 2020. Tesla consistently promised that FSD would be fully capable within a short period of time. But today, almost four years later, FSD still requires a human driver to be in control at all times and we regularly hear about safety concerns and issues in the news.

We bring up the differences between AI products and traditional software products here because we want everyone who has the investment and will to create AI products to be set up for success. If you can anticipate potential hurdles as you're building a novel product, you can better prepare your teams for surmounting them as well.

How are the products similar?

There are a number of similarities between traditional software products and AI products because, fundamentally, AI products actually are traditional software products with a productized AI/ML service built in. For this reason, the similarities we will outline here involve Agile product development as well as data. Native AI products, along with the outputs from the AI pipelines that support them, follow the same building process most traditional software products use. They are also built with a heavy focus on the data that powers them. This is still true for traditional software products.

Agile development

In traditional software development, you follow some methodology to ideate, keep track of your work, and stay consistent with some framework or schedule. Most software companies these days don't use the waterfall methodology anymore and have instead opted to use some version of the Agile, Scrum, Kanban, or Lean methodology. This means most software companies use an iterative and experimental approach to building and shipping products. They take large overarching business goals and translate those goals into specific tasks that will then amount to deliverables by the end of any given sprint. Once they have these deliverables done, they undergo a process of evaluation to make sure they meet varying expectations through testing.

The heart of this approach is *agility*. When you're building out features of your product over time, you have time to test those features both functionally and conceptually. This is an economical way of spending time, energy, and resources on a product or a feature to then see how it's received by your customer base and greater market. The agility this offers is what allows tech companies to be successful: they are able to make changes and adjustments as they build if they see that their product or feature isn't resonating with their audience of users. This will be true whether or not your product supports AI/ML features.

We'd even go a step beyond this and say AI/ML takes the heart of this agility to the next level. Because AI/ML products are consistently building from prior manifestations, they're constantly evolving and adapting to new demands on performance, accuracy, or speed. You can't build an AI product in a vacuum. Over time, AI products will have many transformations, and because of this, they're always in a state of being updated or upgraded to meet the expectations of their outputs. This is particularly true with generative AI products because of the opacity of the deep neural networks that power them.

Since engineers can't really explain why generative models are producing certain outputs, there has to be a continuous process of incorporating changes to the models through external influences, like retraining and bolstering that training with prompt tuning and user feedback. Prompt engineering itself needs to be a permanent fixture in how a company readies a generative AI product for the market because the prompts themselves guide the generation process. Without some way to pivot continuously, particularly in this area, you won't know which prompt strategies are most effective. Agility needs to be a constant.

Data

Then there's the similarity of data. Even if certain software products aren't heavily dealing with your personal data, they are often built as data products in the sense that they are leveraging and storing some data about you or the entire user base to some degree to make certain determinations. Traditional software development will have some feedback loop to a database or be built upon a data pipeline of some kind that passes information back and forth from its UI to some centralized or decentralized repository.

This means that software engineers are working with massive volumes of data in addition to working with source code. We've discussed the data demands AI/ML products have at length over the course of this book, but it's important to note that this is inherently true of most software products out there. Software products are consistently using and accounting for data, whether or not it's a "data" or "AI/ML" product. Acquiring this data from your initial customers if you're launching a new product is going to be the case whether or not you're building an applied AI product.

Note that although this is true, AI/ML products do require a higher degree of data management and data quality practices because the data isn't just used for deterministic qualities of your product; it's also used to train the models powering your product. Feel free to refer back to *Chapters 2* and *3* about ethical considerations of AI products. This is even more important if your AI/ML product is also a generative AI product. Generative AI models require even more data than traditional ML models, so the managing and curating of data at this level is crucial to the productization process when generative AI is involved.

Before we move on to discussing the specific differences in the roles of an AI PM and a non-AI PM, let's quickly summarize the comparisons we've made so far:

Aspect	Traditional software product	AI product
Uncertainty in **return on investment** (ROI)	Lower uncertainty; predictable performance and results after development	Higher uncertainty; ROI depends on model accuracy, data quality, and changing market conditions
Data costs	Lower; minimal ongoing data costs post-development	Higher; data acquisition, cleaning, labeling, and storage are costly and ongoing
Development costs	Lower initial development costs with fewer resource demands	Higher initial costs due to data needs, model training, and infrastructure requirements
Scalability	Easier to scale with standardized platforms and cloud infrastructure	Scalability can be complex due to model retraining, continuous updates, and data dependencies
Profit margins	Typically higher after initial development due to predictable costs	Initially lower due to higher resource consumption and ongoing model maintenance

Development cycles	Fixed, with clear phases (planning, coding, testing, deployment)	Iterative and continuous; involves constant retraining and updating of AI models
Data needs	Minimal; typically does not rely heavily on ongoing data streams	Data-intensive; requires large, high-quality datasets for training and improvement

Table 7.1: Key differences between AI and non-AI products

This table compares traditional software products and AI products against some of the areas we covered in the previous sections of this chapter. What's important to keep in mind with AI products is although they do require upfront time, cost, and resource investment, they have the potential to offer significant returns on the upfront investment once they're optimized and scaled.

How does the role of an AI PM compare with a traditional PM?

As an AI PM, you're often expected to be the person to maintain an intuition about your product. This is true for any PM, but it's especially true for AI PMs due to the depth of technical expertise they need to have in AI capabilities, data pipelines, probabilistic outcomes, model performance metrics, and stakeholder management. Regular PMs still need to develop a sense of product intuition, but their roles are more straightforward. They're often working with more deterministic, fixed outcomes or output. They're also typically interacting with fewer stakeholders who are more focused on business or **user experience** (**UX**) teams.

However, because of the breadth and depth of your knowledge on the product you've managed, as an AI PM, you are often the authority on how your AI product will grow and evolve through the process of building and shipping. Part of that intuition will relate to how you will market and sell your product, what kinds of customer needs and issues your product can anticipate, as well as potential problems that might arise as you start to get into the weeds with building and marketing your AI product. Many PMs might not be aware of the demands AI/ML products will place on them, and this section is primarily aimed at helping PMs build this intuition as they start to navigate the world of AI/ML products.

There are some similarities as well – all PMs will always be tasked with:

- Managing product stakeholders
- Mapping problems to solutions in their products
- Scrutinizing their data analytics and insights
- Communicating with internal and external stakeholders
- Deciding on acceptance criteria that will hold their products' performance
- Accountability, explainability, and ethics of their products.

This is particularly true in the case of generative AI products, where they can generate potentially harmful, biased, and unethical outputs. These areas need to be nurtured whether you are supporting an AI product or not, and it's all part of your evangelism work as a PM.

It's also important to note that the line between traditional software products and AI products is increasingly becoming blurred. This is because most software companies have already started to integrate AI/ML into their existing products or launched AI-native products. PMs that cover a wide variety of products will want to deepen their knowledge of AI/ML as a way to stay competitive within their own fields, whether they plan to go deep into AI or not. Let's turn our attention to how AI/ML products are positioned and built according to their business models particularly B2B and B2C products.

B2B versus B2C – productizing business models

When it comes to building and shipping products, some of the biggest differences between B2B and B2C business models include domain knowledge and the degree of experimentation. In this section, we'll be focusing on these two areas because they have the biggest impact on what productizing looks like for AI/ML products between these two business models. If we expand on the notion that AI/ML products behave more like services, the desired end result of both these business models will be different because they serve different kinds of customers and different overall needs:

- With B2B products, there's a strong need for these products to demonstrate a high degree of *domain knowledge* and a focus on that. Since B2B products often serve a proven business niche, they must often prove they have expertise in this niche.
- With B2C products, we see a focus on *experimentation* because rather than tapping into a business need that already exists, these products are looking to tap into a more collective need that their own customer base may not yet be aware they have. This requires a high degree of experimentation.

In the following sections, we will build on these ideas further.

Domain knowledge for B2B products – understanding the needs of your market

For starters, in B2B products, domain knowledge reigns supreme. This is because the use cases are incredibly specific, and the products themselves solve niche business problems that are specific to certain industries and domains. In order for a PM to be effective in the B2B space, they need to be intimately connected to their customers' needs, workflows, and major pain points because their products solve very specific use cases. This isn't to say PMs in this space have to come from the specific industry they're serving, but they do need to invest a significant amount of time in building empathy with their user base in order to be effective in the following areas:

- **Building credibility**: This time-intensive process can look like a number of things; conducting customer interviews, keeping up with industry trends, and understanding the competitive landscape and the benefits and features their competitors are serving their own customers are all part of this work toward building credibility in the space their product is competing in. This is all preliminary work of understanding the various types of users their product will serve as well. You're setting yourself up for success if you're a PM in this space and you're spending time on establishing clear buyer and user personas and making sure you have a handle on what their individual needs, problems, and "jobs to be done" are.

- **Sales cycles**: The expectation threshold of B2B customers is high because this space is likely riddled with a large number of competing products to choose from. In many cases, these customers are undergoing various rounds of **requests for proposals** (**RFPs**) and in-depth **proof-of-concept** (**PoC**) processes to make sure they are purchasing the right product for their use case. For an AI/ML product, these PoCs can be costly for your organization because you'll need to acquire a large enough sample from your customers, use that data to train your models, and present your product and its capabilities to them once its performance is at an acceptable level to be able to show to your prospective customers.
- **Building within an ecosystem:** This means there are often many eyes on these products, and each of those pairs of eyes may come with its own set of expectations and objections to your product, so you really need to be aware of multiple perspectives when building B2B products. This also means that as you build, you're planning your release schedules with the idea that every release may include features that could impact your customers' individual workflows, so you often have to be mindful of how often and how loaded your release rollouts into production are. This is doubly true because often B2B products need to be integrated into existing enterprise tools, APIs, or other platforms, so you're already working within an established framework of some sort.

This domain knowledge is then built up to the point where you as a PM, and your broader organization as a whole, then become a thought leader in the space you're serving. Because the B2B landscape is a small world, once a product—along with its leadership team—makes a splash, it will trickle across the professional world. Building credibility internally and externally is foundationally proportional to the level of industry expertise that's acquired.

This further reinforces that the expectations of an AI/ML B2B product are quite high. Your users and buyers will scrutinize you on the tech stack that's supporting your AI product as well as its performance and accuracy. They will want you to have evidence of explainability and why your product works, which could make AI products built with generative AI or DL models tricky. This will all be happening in conjunction with the testing and trialing of other AI/ML solutions that are out there to compare them. This means that you will have to be incredibly intentional about your releases to make sure there aren't any lags in performance when you do make your push to production, with the full knowledge that other businesses are relying on your product.

Let's take the example of an AI company that's looking to develop a predictive maintenance tool that will predict when equipment will fail before it occurs in order to avoid downtime in manufacturing environments. The AI engineers who are working on this tool will need to be aware of different types of machinery, as well as the specific failure modes they experience. AI models will need to be aware of the operational limits of certain equipment to interpret sensor data correctly. Domain experts will be needed to identify critical data points like vibration, temperature, or pressure to help the AI product team decide which data is most relevant for predicting failures. In B2B environments, one client's needs can differ drastically from another's, so the product team will need to focus on having a deep understanding of the client companies that will be using their tool and the intricacies and customizations that would be most relevant to their environments.

Expectations are to be managed at every stage of the customer journey. B2B products exist in a massive ecosystem, and companies that use one product might pass outputs from that product to other workflows or their own customers. B2B products frequently integrate with other software solutions to streamline workflows or enhance operations. For instance, a project management tool might take data from a time-tracking application to generate reports. Data from those reports can be fed to a team collaboration platform to make sure stakeholders stay informed, and so forth.

This places an expectation on B2B companies to maintain their company's and product's reputation and be more transparent about their marketing efforts. Long-term relationship building is foundational to the success of a B2B product because the contracts are often multi-year and sometimes the sales cycle is too. This introduces a lot more strain on AI/ML product teams to shy away from black-box algorithms and sell products that can stand by their outputs because there's a greater need for transparency, accountability, and explainability. If a B2B product fails to deliver on performance or expected results, businesses want to understand why, often in the form of a formal live meeting. B2B company customers have a tremendous amount of leverage because the stakes are high all around, as are the contract costs. B2B products benefit from having engaged customers who want to see a certain level of performance from their products and that feedback and collaboration help improve products.

Experimentation with B2C products – discover the needs of your collective

With B2C products, because there isn't such a hyperfocus on the domain you're serving, the pressure on the market you're serving is a bit more relaxed, but this creates another kind of pressure: the pressure to build a product that appeals to a much wider audience universally. Casting a wider net brings other areas into focus, and the following are some key considerations:

- **Gaining customer insights**: As a PM building in a B2C environment, you're going to approach your product and the market it serves more experimentally and derive insights about what's most useful to your customers by tracking how they behave with and use your product. You can conduct focus groups, nurture beta testers, or interview your customers through in-product surveys, but because you can't conduct customer interviews in the way you might for a B2B product, you're left with understanding your customers' impressions of your product through their in-product behaviors.
- **Enhancing customer experience**: This business model also consolidates all your user and buyer personas in one because, typically, the person who's using your product is also the person buying it. Understanding the main drivers, desires, hopes, and dreams of your customer base becomes a very nebulous task because what you're trying to capture are underlying needs, pain points, and moments of delight that will apply to all your users at once. This complicates a PM's ability to empathize with their end users. Because most of these users aren't signing the kind of year-long contracts you see with B2B products, the pressure to keep these consumers charmed is constant because they can leave at any point. This puts pressure on product, leadership, and development teams to consistently provide their customer base reasons to stay and choose them. Creating a streamlined, intuitive, standardized experience is one thing but deciding on which experience to choose above all others is another.

- **Investing in data analytics**: Because of the bird's-eye view B2C products enforce on their builders, this means PMs have to be very discerning with their data analytics and metrics. Investing in understanding their customer lifetime value and customer acquisition costs and tracking those metrics is an important part of staying profitable and sustainable in the B2C landscape. B2C offers PMs an easy outlet for applying AI/ML toward the acquisition and retention of customers since they have such few touchpoints with the end users of their product. PMs in the consumer space also need to hone in on the demographic and individualistic qualities of their consumers to better understand what to build. If you see that it's mostly people within a geographic area, gender, generation, or subgenre that appeal most to your product, you might start to build future features and releases in your roadmap with them in mind.

We've mentioned many times that AI/ML products are experimental in nature because you want to leverage AI/ML in ways that will impact your product most obviously for it to be worth the top-heavy investment it requires. This is doubly true for B2C products because you're building and using AI to deliver something that saves your consumers money or delights them, as well as using the data your product produces to decide on how to pivot your product. B2C PMs are reliant on data and analytics to make global decisions about their products on a regular basis. Although this is changing for the most part, traditionally in the B2B space, your marketing efforts are most oriented toward your buyers because they're the ultimate decision-makers. This is in sharp contrast to the B2C space because your marketing efforts are directly linked to collective information you can derive and infer from your data.

Let's consider an AI company that wants to offer a personalized styling tool that will suggest clothing to customers based on their preferences, purchase history, and feedback. Because the tool will be applicable to all users at once, experimenting through fine-tuning performance and optimizing UX will be the key to success. This might require experimenting with different algorithms, learning paradigms, or filtering methods to find the best combination that will capture the preferences and behaviors of their user base the best. A/B testing and iterating on personalization criteria will allow product teams to try different recommendation formats, styles, and explanations for users and customers that will maximize purchases, customer satisfaction, or positive feedback. Then, the most effective approach can be discovered and scaled across all users.

Experimentation is fun to a point, but at the end of the day, it has to deliver. By far the greatest pressure this places on AI/ML consumer products is the pressure to perform well, maintain their enormous user base, and keep their consumers happy. This means that consumers want to use an app that does what it says it's going to do in a way that's visually appealing. Because they aren't as concerned with the downstream risks of using your product as B2B customers are, consumers don't particularly care why it's working, just that it does work. This means that the issue of explainability is minimized in the B2C space, and the use of black-box or DL algorithms is less scrutinized.

B2B and B2C business models come with their own blend of challenges, but at the end of the day, players in both business models must understand their customers enough to create products that actually bring them value. Once you've built something of value for your customers, you enter a new phase. This next phase is about delivering that value consistently enough to not only win customers but keep them in the long run. In the following section, we will take a look under the hood to understand the necessary elements of delivering value consistently with the help of **MLOps** or **AIOps**. Both will be used interchangeably in this book.

Using AIOps/MLOps

Maintaining trust, reliance, and consistency within your internal product teams as well as with your customer base involves an act of committing to specific rituals. Ritualizing the acquisition of clean data, tracking the flow through your infrastructure, tracking your model training, versions, and experiments, setting up a deployment schedule, and monitoring pipelines that get pushed to production are all part of the work that needs to be done to make sure you are keeping on top of the comings and goings of your AI/ML pipeline. This ritualizing is what's referred to as MLOps or AIOps. In the following sections, we will explore a few important considerations that relate to AIOps and MLOps.

Consistency and AIOps/MLOps – reliance and trust

MLOps' greatest contribution to your business is maintaining the consistency needed to build an AI/ML product that lasts. Let's explore the benefits of AIOps/MLOps and how they help you stay consistent:

- **Managing ML pipelines**: If you're managing an ML pipeline, you will need to learn how to depend on an MLOps team and set up your team for success. You don't want to get caught losing $20,000 in 10 minutes (as we saw in our *Profit margins* section earlier in this chapter) and have no leads for where the problem stems from. At the very least, you should have some idea of where the problem is stemming from. MLOps is able to help with creating and managing the ML pipelines themselves, scaling those pipelines, and moving sensitive data at scale. Ultimately, the risks of compromising your customers' reliance on and trust in your product are great.
- **Productizing AI/ML**: We've spoken at length in this chapter about productizing and what that looks like in different contexts of AI product management, but MLOps is actually where the productizing functionally happens. Taking a service and splitting that service apart into smaller pieces that are managed and standardized into reproducible and regulated segments is the work of "productizing," and in that vein, MLOps is really the vehicle we use to truly productize AI/ML. In order to build the consistency and credibility customers can expect from your product, the ritual of MLOps needs to be cemented into your process. You never know when you'll need to revert to a specific version of your models or zero in on an experiment that had the right mix of factors. MLOps creates the organization and focus your team needs to have in order to have a handle on such a wide variety of experiments.

Tools like Splunk, Dynatrace, and Moogsoft are used for AIOps to provide analytics and monitoring tools that help IT teams analyze machine data and logs. These are AI-powered monitoring solutions for applications and infrastructures used to keep your AI program working properly. In many cases, they also have proactive monitoring features as well as incident response and management features to help troubleshoot when (and for how long) something is going wrong.

Tools like MLflow, Kubeflow, Data Version Control (DVC), Seldon, and Weights & Biases are used for MLOps to manage AI/ML lifecycles from experimentation to deployment. Managing and scaling ML workflows, maintaining versions of AI/ML projects, A/B testing, managing different deployment strategies, and visualizing performance metrics are all functions these tools can offer to ML, engineering, and data science teams.

Now that we've covered some of the benefits of creating an MLOps organization to help keep track of your AI/ML pipeline, let's explore how to build on that pipeline and evaluate the state of the models used. In the following section, we will reiterate some concepts around testing, retraining, and hyperparameter tuning to ensure your AI/ML pipelines are routinely being refreshed and optimized for performance.

Performance evaluation – testing, retraining, and hyperparameter tuning

Without having a built-out AI/ML pipeline that validates, trains, and retrains regularly, you won't have a great handle on your product's performance. Establishing a baseline, as well as specific benchmarks, for trust, reliance, and consistency in AI/ML performance evaluation is essential for building effective AI/ML operations. AI PMs can expect to help build standards and benchmarks around this. How you define your standards is up to you, but the most important thing is that you do define it somehow. The following are some typical MLOps standards you can potentially apply within your AI/ML program:

- **Accuracy**: Maintain an accuracy threshold that's balanced with precision and recall to make sure your model isn't favoring false positives or false negatives. You could also set a standard for model performance to ensure that 95% of your model predictions are within an acceptable margin of error over a defined period.
- **Latency**: Establish latency thresholds for real-time applications so that you're getting responses from your models within an appropriate time if your AI product is trying to detect fraud within a certain window of time, for example.
- **Audit trails**: Keep a record of model versions and performance evaluations.
- **Trust metrics**: You could employ user surveys to establish a trust score so that over 80% of your users feel confident about your model or product's predictions.
- **Variance**: Monitor your model performance so that there's no more than 5% variance on key metrics, like accuracy or recall, over a defined period to prove your product is consistent.

MLOps helps us with accentuating the importance of retraining and hyperparameter tuning our models to deliver performance. Here are some tasks that your MLOps team will likely be performing:

- Your MLOps team will essentially be made up of data scientists and ML and DL engineers who will be tasked with making adjustments to the hyperparameters of your model builds, testing those models, and retraining them when needed. This will need to be done in conjunction with managing the data needed to feed this testing, along with the code base for your product's interface as well.
- In addition to testing and validating the models and working to clean and explore the data, MLOps team members also traditionally do software testing, such as code testing, unit testing, and integration testing.
- In many cases, your AI products will effectively be traditional software products that incorporate AI/ML features in a subset of a greater ecosystem that's in line with traditional software development. This means that MLOps may, in many cases, ensure your greater product is functioning along with the AI/ML deliverables and outputs. This will be true for generative AI models as well, which will need to be regularly evaluated for diversity, coherence, perplexity, and factual accuracy. Maintaining standards in performance is especially important in cases where you have a lot of competition as well.
- Another major area MLOps is well suited for is minimizing the risk of data drift and system degradation. We covered a few different types of drift in earlier chapters of this book, but we'd like to reiterate here that this is a risk that can sneak up on you. Model degradation can happen for a number of reasons across all AI/ML models, including generative models. Perhaps there are differences between assumptions that are made with the data in training and production. Perhaps there's been a change to the data itself. Perhaps there are unseen biases in your training data that were never picked up on. Whatever the reason, continuous monitoring of models in production by MLOps will be your best defense in picking up on these nuances and changes in AI/ML outputs so that the risks from any number of issues within the ML pipeline are minimized as much as possible.

> In *Chapters 1* and *2* of this book, we covered the concept of continuous maintenance, which consists of **continuous integration** (**CI**), **continuous deployment** (**CD**), and **continuous training** (**CT**). These basic areas of MLOps are mirrored in DevOps, which is common to traditional software development. The main differences are:
>
> - **CI**: In MLOps, CI isn't just about testing and validating code—it's also about validating and testing the models, data schemas, and the data samples themselves.
> - **CD**: CD isn't just about deploying a software package in MLOps but about nurturing an automated ML pipeline deployment process that's optimized for deploying the model prediction service or for automatically scaling back to an earlier version if there is trouble ahead.
> - **CT**: In MLOps, CT isn't just about testing software packages themselves but about retraining and testing the models that are actively being relied upon.

In this section, we've built on the idea of building consistency in how we manage our AI/ML pipelines and reinforced the importance of maintaining high standards in the performance of our AI/ML pipelines. This shouldn't be viewed as a nice-to-have but rather a need-to-have. The performance and quality of AI/ML models can suffer for many reasons, so this consistent practice of scrutinizing performance is meant to ensure your product's performance doesn't come at the risk of the trust you've built with your developers and customers. In the following section, we'll be discussing the importance of maintaining strong relationships, whether they are internal or external.

Feedback loop – relationship building

Continuously monitoring and reinforcing the legitimacy of a complicated system, such as an AI/ML pipeline, is all in service of the ultimate goal of building relationships that last. Relationships between your company and customers, your development team and your sales team, and your MLOps team and your leadership team are all forged through this work of building and going to market with your AI/ML-native product. In AI/ML products, the feedback loop is everything. Nurturing a strong relationship with the builders of these products and the customers they serve is the underlying work of the PM. Many layers of work go into productizing, but at its most basic level, this work is really just an elaborate feedback loop. The feedback is important for discovering areas for improvement and informing future retraining and fine-tuning.

We haven't discussed marketing much in this chapter, but this will also be an integral part of maintaining this feedback loop. Finding the right words to use to describe your product and reach your audience (also known as product language fit) will be a big part of productizing. You will have to create marketing collateral, advertisements, and sales scripts that will all convey the value the product you're building with AI/ML will have for your customers and end users. Those communications will need to be reinforced by the feedback from the customers you already have who are finding value in your product. Building product collateral and expressing that through your various marketing and sales channels will go a long way toward level-setting expectations with your customers. You can go back and refer to *Figure 7.1*, which illustrates what this feedback loop typically looks like.

As we've seen earlier in this section, customers come with expectations, whether they're business users or consumers, and it's the task of the AI/ML PM to take those expectations and deliver something that aligns with them. Understanding the risk and promise of AI/ML and how it compares and contrasts to traditional software development, understanding the challenges and opportunities in your business model, translating all that to reproducible, repeatable tasks internally, and demonstrating consistency with your product's performance in a way that aligns with your customers' expectations are all involved in productizing the AI-native product. You'll know you've successfully done this when you have a loyal customer base that wouldn't dream of parting with you.

Case study

Building off our case study example from *Chapter 6*, in this chapter, we will be deepening our understanding of Waterbear and their flagship B2C product, Akeira, to better understand some of the principles outlined in this chapter and how they may appear in our working example. We will use the flowchart shown in *Figure 7.1* to inform how we break down this chapter's principles in our case study example:

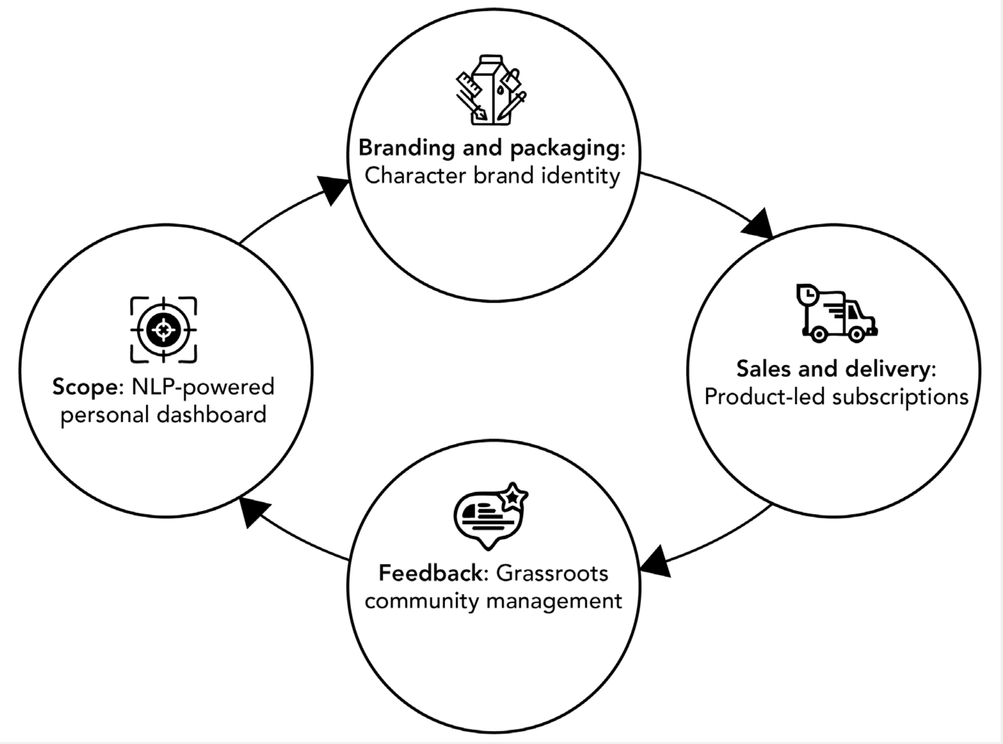

Figure 7.2: Productizing basics for Akeira

Here are the steps involved:

- **Scope**: Here, the product is a B2C mobile application that uses **natural language processing** (**NLP**) to digest the data users feed into a dashboard that is used to track progress toward their established goals. End users can expect that the dashboard will be refreshed regularly, every time they make a new submission into the app. Waterbear currently offers Akeira to women over the age of 18, but it's in the process of creating different use cases for girls under 18, boys under 18, men over 18, non-binary people under 18, and non-binary people over 18. Each will be a different product. Currently, their language models are training on each individual customer's data, although they are testing a hybrid model that will merge smaller localized language models with training data from their entire customer pool.

- **Branding and packaging:** Although Waterbear does have one integrated blog, they plan to have separate brand identities based on the characters of their various products. Akeira is their flagship product and she is represented as a woman to mirror her intended user base. Because users are sharing personal information about their lives and emotional states with the app, Waterbear wanted to personalize the experience by building a personality and vibe around who Akeira is. Many of the end users of Akeira are Gen Z and Millennial women, so the feel, look, and sound of Akeira are palatable to that user group.
- **Sales and delivery**: With Akeira and Waterbear, the delivery is through the app store on all major device carriers. Because it's a consumer product that functions in a product-led organization, most of the sales come through marketing efforts, easy installation, ease of use, and customer satisfaction-based renewals. Akeira is powered through a subscription, which consumers pay a monthly fee for.
- **Feedback:** Currently, Akeira has over 200,000 monthly active users and the only sustainable way for its product teams to manage customer feedback is predominately through in-app surveys. Their team of customer success specialists also responds to feedback online through social media sites, such as Reddit, TikTok, Instagram, Facebook, and X (formerly Twitter). Occasionally, they host pop-up events in major cities. All feedback is also fed through Waterbear's NLP capabilities so that areas of continuous improvement and new features can be discovered. This is how they discovered new use cases and cohorts. Based on the feedback from their existing users and potential users who were not yet eligible, they discovered that there are significant user groups that would benefit immensely from their product.

Summary

In this chapter, we learned that the act of productizing involves taking a concept, a service, or a piece of technology and developing it into a commercial product that's suitable for the customers you're looking to attract. As we've seen throughout the chapter, this work isn't just a matter of getting your product up and running and creating a landing page for your potential customers to magically find. Productizing involves critically understanding the business model you're working in and the ultimate audience you're building for. Remember that AI products can be thought of as AI/ML services that are being built into traditional software products. This means that another big part of productizing for AI products involves the standardization and ritualization of the AI/ML service in a way that's repeatable and predictable for the internal operations teams as well as the customers who will come to rely on your product.

If you're able to understand your market, build internal structures to make sure there's consistency with the outputs of your AI/ML pipelines, and communicate that consistency through your marketing efforts, as well as the ongoing performance of your product, you've successfully productized. But productizing may not be enough. Depending on the specifics of the market you're serving, you might have to customize your product even further for specific use cases, verticals, customer segments, and other peer groups. Because AI/ML model performance is so dependent on training data, there might be collections of data that perform differently when run against one model. Further specialization might be in order. If you find this applies to your product and market, read on to *Chapter 8*, where we will build on the concept of customization.

References

- *Past the AI hype: When is AI actually profitable*?: `https://www.forbes.com/sites/forbesfinancecouncil/2020/09/01/past-the-ai-hype-when-is-ai-actually-profitable/`
- *How to scale AI in your organization*: `https://hbr.org/2022/03/how-to-scale-ai-in-your-organization`
- *SaaS metrics 101: Gross margin and net profit*: `https://www.capchase.com/blog/saas-metrics-101-gross-margin-and-net-profit`
- *Guide To to data cleaning: Definition, benefits, components, and how to clean your data*: `https://www.tableau.com/learn/articles/what-is-data-cleaning`
- *Why commercial artificial intelligence products do not scale*: `https://ai-forum.com/opinion/why-commercial-artificial-intelligence-products-do-not-scale/`

8

Customization for Verticals, Customers, and Peer Groups

In this chapter, we'll understand how products evolve and segment themselves across **verticals**, **customers**, and **peer groups** (customers of similar sizes that may not necessarily be in the same vertical but are complementary). The purpose of this analysis is to understand how you can start to think of product management's role in orienting AI products for specialized groups and, ultimately, what AI allows you to do for those groups. AI and ML are powerful tools but have a general power that needs to be applied specifically in order for their value to be fully appreciated. Adjusting AI models to meet industry-specific needs, personalizing AI solutions based on customer demographics, behaviors, or company size, and leveraging data to share insights across industries and sectors are all ways to create customization when bringing forward solutions. In many ways, the work of a **product manager** (**PM**) is to make that value as obvious as possible to everyone, from customers to developers. For this reason, we will be discussing the various ways products can be contextualized in terms of domains, verticals, and use cases.

As we've discussed in previous chapters, the role of a PM incorporates a lot; it's multifaceted in nature. You're involved with designing the product, organizing the workflow as that product evolves, analyzing feedback from your customers, incorporating that feedback into your overarching business goals, researching new methods and improvements, building the greater product strategy that aligns with your company's strategy, and, ultimately, communicating across all levels to all your stakeholders, developers, and leaders. This requires a lot of understanding, intimacy, and commitment to your product and its success!

When you incorporate the role of a PM into an AI PM, you've got a few more areas to consider. You'll want to understand the options you have in terms of ML algorithms and data models, as well as understanding how AI will move the needle as far as your acceptance criteria and key metrics are concerned. Particularly now, with the accessibility of generative AI models and APIs, you'll also want to try to understand how generative outputs can be used to leverage the value and strength of your product. You'll want to have an understanding of data privacy and AI ethics to make sure you aren't building something that might cause harm to any of your customers or the business at large down the line.

Above all, the role of an AI PM will require a level of data fluency and literacy. Gaining an intuition about how to use data, how to apply data to specific use cases, and understanding a bit about statistics and the data models themselves will be crucial to making sure you've got enough of an understanding to be confident as an AI PM. This area is, in many ways, a prerequisite to being able to start your AI PM career. As an AI PM, you'll be working with data heavily. Critical thinking and complex problem-solving will be necessary to build a successful product that best aligns with the specific use case for your AI product. To help you develop these skills, we will explore the most common use cases, domains, verticals, and peer groups that AI is leveraged for.

The topics we will cover in this chapter are as follows:

- Understanding domain spaces and how AI can be optimized for these domains
- Fluency with various verticals that are seeing a high saturation of AI products
- Understanding user behavior analytics
- Learning from peers in your space through thought leadership and consultancy

Domains – orienting AI toward specific areas

When discussing **domains**, there are really two major domains you'll want to invest significant time establishing credibility in. The first is in *understanding AI concepts* themselves, which we've already covered extensively throughout this book so far. The second is in *understanding how AI is helping certain domains achieve success*. Considering that the part of the book we're in is focused on building an AI-native product, we will focus on how an AI PM can achieve this as they're setting out to build a new AI product.

Depending on the industry you're in, you're going to want to understand your space well enough that you have an understanding of who your direct and indirect competitors are. This might be more straightforward or not, depending on the space you're working and competing in. Gartner has a great tool called **Magic Quadrant**. The four areas of the quadrant are as follows:

- **Leaders**: Established players in a market that have been most reliable and prevalent. They demonstrate the most complete vision as well as the highest ability to execute that vision.
- **Challengers**: Newcomers that are threatening the dominance of leaders in a given market. Their completeness of vision isn't quite where it is for leaders, but they are able to compete on execution.
- **Visionaries**: This group has a lot of vision but their ability to execute it is limited, so they're competing with leaders but they're not able to execute quite as quickly or comprehensively.
- **Niche players**: This group is able to focus on a limited amount of use cases and serve that select group quite well. They're not expanding their product suite to meet a higher vision, nor are they aggressively trying to gain a higher market share. They're happy in their niche.

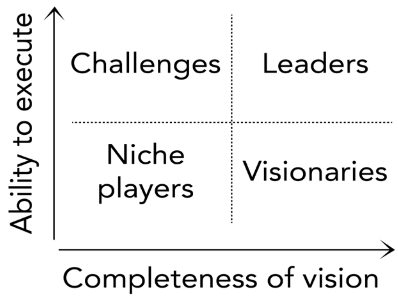

Figure 8.1: Adapted from the Gartner Magic Quadrant

As an AI PM building a novel AI product in your space, you'll want to have an understanding of all four kinds of competitors in your chosen space or domain. The four quadrants themselves are built on an x axis based on the completeness of vision, and a y axis based on the ability to execute it. We encourage you to use Gartner's Magic Quadrant as a resource when evaluating your organization's placement in your own market, as well as the placement of your biggest competitors. Your ability to gauge where you fall along these fault lines will best prepare you for coming to market with a new AI product, particularly now that so many AI products are still in their early days.

When building the AI-native product, a lot of initial work needs to go into ideating and positioning your product because it's so new and it's entering a new market. Given the cost and investment of getting an AI product up and running, you'll want to be deliberate when building that investment. In the following sections, we will be focusing on the work of positioning a product appropriately for your market, understanding how your product design will actually serve that market, and how to build an AI product strategy that will make the most sense for your market.

Understanding your market

For any specific domain, Gartner's Magic Quadrant helps companies identify and orient themselves against their competitors in terms of execution and vision. Doing competitive analysis to understand where you are and where you perceive others to be in terms of your products and AI capabilities today will help you determine your future strategy for your product or suite of products. Let's look at what the different categories are:

- **Leaders**: The leaders demonstrate a proven history of executing the most complete vision. They can show you what's been working and what needs the main market players in the domain are predominately addressing. This will be valuable for you to get a sense of what the addressable market actually wants out of the companies that serve them. Leaders are also great for showing their domains what a sustainable execution of domain pain points looks like, as they've been around for a long time and fellow PMs can learn from the products they've released and how the greater market has received them over the years. Leaders offer a great foundation for research when you're in a new domain as a PM. In this section, you'll find companies like Microsoft (Azure AI), Google Cloud (AI and ML), and IBM Watson, which offer comprehensive AI tools, widespread adoption and market dominance, innovation, and scalable solutions that have stood the test of time.

- **Challengers**: Though challengers in the quadrant show less of a complete vision, they're able to address the needs of their market well enough to take some of the market share away from leaders. This is a helpful dynamic to understand when you're thinking about your own product. If challengers are able to take market share from the leaders in their domains, it means there are potential fault lines to tap into. Those fault lines are areas where leaders are not addressing the needs of their market adequately and the most successful challengers are able to find opportunities in those fault lines.

  In this section, you'll find companies like AWS, Oracle, and Salesforce, which offer high levels of execution in AI/ML platforms but aren't leading the charge in terms of innovation.

- **Niche players**: Niche players are not overachieving in the completeness of their vision nor their ability to execute it, but it's exactly what gives them their superpower, as the niche players are able to show you where the edge cases might be. This could help you potentially identify features or use cases for your product as you're building and designing that have some proven success. These edge cases might be so niche that the challengers and leaders themselves don't find them strategic to build into their existing product suites, but for an emerging product, this knowledge could be valuable. Getting an understanding of what niches exist within your domain will show you where there might be latent areas of growth that are currently untapped or underserved. In this section, you'll find companies like H20.ai, SAS, and Cloudera that focus on specific verticals and use cases.
- **Visionaries**: These are the players that have a robust vision for their product suite and, likely, their entire domain. Visionaries are able to show you what's possible from a product in your chosen domain and will likely include some of the most cutting-edge features you're likely to find in your space. As an AI PM, you'll see most of your direct competitors here because AI features are still slow to adopt in most domains, and we would wager that most of these players are starting to experiment with AI capabilities. In this section, you'll find companies like DataRobot, UiPath, and Palantir, which offer strong visions but in more limited scopes.

If you're getting started in AI in your domain, and you have the privilege of being involved with your product strategy from the very beginning, understanding the emerging visionaries in your domain will be a great place to start. Today's visionaries may be tomorrow's leaders, so you'll want to get cozy with the names in this category for your chosen domain.

You can do your own market research, but either way, you're going to want to have some sense of the landscape that makes up your domain so that you understand your competition, the problems they're trying to solve, and the customers that are in your domain as completely as possible. Working with your marketing team or building a comprehensive understanding of your competitive landscape is a crucial layer to building your AI-native product. In order to build something of value, you have to build something that addresses the deepest customer pain points you're trying to solve with AI. You can use the following checklist to determine who you are today and who you want to be in your market tomorrow:

- **Identify key players**: Map out your key competitors against yourself in the four quadrants to understand strengths and weaknesses. What are the problems your competitors are trying to solve? Who are the customers that make up your target market?

- **Research current trends**: Analyze the impact of emerging technologies and trends to understand what influences are having the biggest impact on your overall market. What are the gaps in the existing tools out there?
- **Evaluate your positioning**: Review the quadrants and the players in your market that are in them to determine if your current product and company strategy align with what you're seeing. If it doesn't, think about how you might get to where you want to be.
- **Inventory capabilities and strengths**: Are there products or features that are most prominent for the companies in your market? Are these relevant to where you want to take your company?
- **Customer feedback analysis**: Are there gaps or pain points in some of the perceived strengths of your competitors based on the quadrants? How does your product stack up?
- **Analyze market share**: For your most prominent competitors, what does their market share and growth trajectory look like?
- **Market segmentation**: Are there certain segments within a market that your competitors are servicing most? Are these competitive or complementary to your company's focus?
- **Feature level**: At the product level, compare the features, capabilities, and pricing of your products across competitors in all four quadrants. Do you notice adjustments you can make to your roadmap?
- **Strategic partners**: Can you identify potential partnerships and collaborators within the four quadrants? Are there complementary partners that you might collaborate with?
- **Find your disruptors**: Which companies within the quadrants may pose the greatest risk to your company or product? Are there elements you can introduce that could limit this risk?

Beyond the pain points themselves, your product will also begin to collect customer data. This means that you're not just building a product that addresses your customers' problems; you're also improving how your product functions with the data your product is collecting from your customers. We will talk about this specific point further in the UEBA section, but for now, keep in mind that understanding how much and what kind of customer data to collect, in addition to how to use that customer data, will be part of the work you need to do while doing the market research to fully understand the market you're looking to serve.

For many AI PMs, you likely already have some familiarity and credibility built in your domain. It is possible to grow your reputation in your chosen domain by bringing something novel to your company, organization, or industry. AI offers a vehicle to do this but, again, you first must have a clear answer to why you want to be using AI that goes beyond the marketing fluff so that you can stand out from the competition. Making sure that you're understanding your market needs, understanding how your product will specifically address those needs, and building and communicating a product strategy that aligns all that into one path forward will be a great way to make sure AI is properly oriented toward your chosen domain.

Again, this will be a prerequisite for ideating on the early stages of your product design and strategy. Feel free to refer to *Chapter 2*, where we discuss **new product development** (**NPD**) stages, and *Chapter 6*, where we expand on those ideas for AI product development stages.

Understanding how your product design will serve your market

After you feel like you understand the domain, the major players, the way they're leveraging AI, and some of the main customer pain points you're looking to solve in your chosen domain, you're going to be ready to start to ideate on your product. This will come with a whole slew of activities. Before all the wireframing, building version one of your product strategy, and coming up with a roadmap, you're going to want to be really clear on why you're using AI to address this market. Building an AI product to address your market without having a clear understanding of why it takes AI to solve this problem is an easy way to waste the money you worked hard to raise.

Yes, there's a market splash and many start-ups embrace AI just for the funding, but this alone will not have enough substance to get you through the stages of responsibly building an AI product that will perform well in your market. Before you go out and build an amazing product that solves all your customers' problems with the power of AI, you will need to solve the fundamental question of *Why AI?*

Once you have a clear idea of *why* you're using AI to solve your customers' issues, you have the building blocks of your marketing message. Because we're at an integral time for AI products, just coming to market with AI features is enough to put you in the position of having to evangelize AI in your domain. Sure, some industries and domains are seeing lots of AI features already, as we will cover in the next section, but by and large, you'll likely be among the first, so this question of *Why AI* will keep coming up from your customers, from your competitors who insist their product is superior, or from your design and building team internally.

As you begin to ideate, you will need to keep this answer, along with your customers' major problems, at the front of your mind so that you remain customer-focused and dedicated to your message. As you build through the initial versions of your product, always keep asking how AI is helping you solve your customers' most frustrating and integral needs. As a PM, you're an energy-generating force within your company. This means that your work of ideating and building a solution that works for everyone is also meant to stimulate your leadership team. Keeping your focus on *Why AI*, as well as the primary problems of the customers you're serving, means you're keeping your leadership team focused as well.

The next decade will show us new ways of applying AI in a way that addresses specific domains. If a great solution doesn't yet exist, see whether you can dream it up yourself. If you can pull it off, your domain will reward you with thought leadership bragging rights, an ever-growing customer roster, and the knowledge that you were able to forge something memorable in the wilderness. Many still think of applied AI as the Wild West because technologists know that by 2030, we will be looking back on some amazing breakthroughs and use cases of AI that we haven't even thought of yet. It's important to keep this spirit of curiosity and creativity alive in such a nascent field.

This applies no matter which domain you're in. Some domains will be more advanced as far as AI adoption is concerned and this will make your work as an AI PM more and less difficult because each level of advancement of AI has its own slew of challenges and opportunities, but the process will be the same regardless. If you're in a domain that is more advanced as far as AI adoption is concerned, you'll be more interested in which AI features your competitors are using to address certain needs and decide whether that's a direction you want to go in as you refine your product strategy. If you're in a domain that's less advanced, you'll have to look to other industries for inspiration for AI adoption and you'll be able to make a more eye-catching splash with your own product in your market.

Either way, you will be able to successfully reach a market when you're able to apply all you've learned in a way that offers your domain something that's currently missing from the offerings that already exist. As an AI PM, you don't need to be an expert in all areas of AI but you do need to have enough of an understanding of your market and AI capabilities/strategies/algorithms that will best serve the exact use cases you're trying to apply AI.

Keeping your leadership team, **go-to-market** (**GTM**) marketing/business development/sales team, and development team aligned with your customers and the AI solutions you're bringing to the table means you're helping your entire stakeholder team focus on your main goal, which is to build an AI product that best addresses your market's needs. This seems like an obvious point, but as many PMs know, once you get started with the day-to-day work of building and evolving a product through your various sprints, losing sight of your main goal is easier than it seems. There are so many specialized skills that an AI PM needs, so keeping this focus is important when you start getting into the weeds with your design, choosing which AI algorithms to go with, and so on. The following is a list of considerations and questions to get you started:

- What specific customer pain points are you aiming to solve?
- What is the impact these pain points currently have on your customers' experience?
- Why is AI the right solution for these pain points?
- What unique capabilities does AI bring to the table that traditional methods don't?
- How are the key competitors in your market currently using AI to address similar problems?
- What gaps exist in their offering that your product can fill?
- How will you educate and explain the benefits of your AI solution to customers and stakeholders?
- What potential challenges could arise from using AI in your product?
- Is there a plan for how you will address concerns related to data privacy, model bias, and ethical applications of AI?
- Are there strategic AI milestones in your product?
- How will your AI product empower the overall business goals at your company?
- Which AI algorithms and models will you choose and why?
- How are these choices supporting your goals of solving customer problems?
- How can you make sure your focus remains on these main goals throughout the product development process?
- How will you limit scope creep and distractions during sprints?

> For more discussions on design, feel free to jump to *Chapter 9*, where we will be expanding more deeply on product design for the AI native product, and *Chapter 15*, where we will discuss the role of design when traditional software products evolve with AI features and capabilities.

Building your AI product strategy

Once you've done the work on designing a product that addresses your customers' and the market's major pain points you want to solve with AI, you're then in a position to start to build out your overarching product strategy. Obviously, your product MVP will start somewhere but after that, there needs to be a course for your product vision to grow toward. Remember that completeness of vision was the x axis in Gartner's Magic Quadrant? The reason for this is that an addressable market has a variety of needs, ranging in complexity. Your product MVP will only solve a small fraction of these needs. Your leadership and company goals will likely not be able to solve all the needs of your market, but you will need to work with your leadership team to define what you're looking to do for your market at large, as this will directly inform your product strategy.

Remember, building a product strategy falls on the PM but creating company goals falls on leadership. You can consider your leadership team as offering you the crucial outputs (the company goals) you need to build a robust product strategy, vision, and roadmap that aligns with your company's business objectives. This is an important distinction to have. PMs can put a lot of pressure on themselves and it's something I wrongly internalized early on in my own PM journey as well. Being a PM doesn't mean all the grand ideas have to come from you. Lean on your leadership team to inform your work.

The work they should be doing is defining the company's overarching vision, mission, and objectives. Without these clearly defined, it will be difficult for you to formulate any real product strategy because there won't be anything to base it on. Once you can take these overarching company goals and turn them into a product vision, strategy, and roadmap, you're able to communicate this to your development team and start the work of building the actual product, which will result in an MVP all the stakeholders involved can be proud of that satisfies the minimum solution needed for your company to go to market with your product.

In order to build a product strategy and vision that can grow with your market, you will need to see how your market has been changing and have as complete an understanding as you possibly can to look at where this market is headed. Countless articles come out every day about the trends for particular industries and domains. As an AI PM, it's even more important for you to keep ahead of these articles so that you have a clear picture of how your market is changing and evolving.

Many of your internal stakeholders, along with your customers, will be looking to you for thought leadership and communications about how your company or product suite is growing with the market. Staying aware of these market trends, particularly as they align with your AI strategy, will be necessary for you to keep maturing your product strategy responsibly in a robust, evolving, unpredictable world.

Nowadays, specific domains can have hundreds or thousands of competitors in the space. Particularly now that so many start-ups can erupt and work remotely first, we're seeing a growing number of companies coming into virtually every domain. Being able to communicate your company's offering in a way that meets your customers where they are, using the technical language or jargon they're familiar with, and being able to show them why your product is the best choice through AI will be your challenge and opportunity as an AI PM.

While we will be going over some of the common use cases as we continue with this chapter, we do want to stress one important point before we complete this section: AI is incredibly customizable. Its basic building blocks are data, and what you choose to do with that data is really what gives AI the power it possesses. If you're an eager, talented PM in the AI space and you're passionate about solving your customers' major problems through a novel product you're helping build, you can build something highly customizable to your customers' specific needs.

The following is a list of considerations and questions to get you started:

- What are the overarching company goals set by leadership?
- How do these goals translate to specific product objectives?
- What is the long-term vision for the product?
- How does this vision align with company goals and market needs?
- How have the key pain points in your market evolved over time?
- What trends are emerging?
- How are your primary competitors positioning their AI products?
- What AI features or capabilities differentiate your product from theirs?
- What specific needs do your target customers have?
- Are there other segments in your customer base for more tailored solutions?
- What recent articles or reports can you use to source trends in your industry?
- Which will you be checking regularly to maintain a pulse on your market?
- How can insights from these reports inform your product strategy?
- How will you communicate your product vision and benefits to stakeholders?
- What messaging will resonate best with your target customers or segmented customers?
- How will you make sure there is alignment between the product marketing and development teams?
- Which recurring meetings, channels, or methods will you set up to make sure communication lines stay open?
- What is your go-to-market strategy for launching the product?
- How can you encourage creativity and the exploration of new ideas and innovations within your team?

> For deeper discussions about product strategy, feel free to refer to *Chapter 10*, where we expand on performance benchmarking and metrics that align with your product strategy. You can also jump ahead to *Chapter 14*, where we discuss product strategy in the context of evolving traditional software products into AI products.

Now that we've spent some time seeing how to best position a product for your market and the considerations of building an AI strategy that serves the customers you're trying to attract, we can take a look at some of the most popular verticals for AI. Each vertical will have a series of popular use cases that we see time and again, and this can shed light on some of the greatest benefits AI can give certain verticals.

Verticals – examination of some key domains

We discussed general domains in understanding how AI is to be oriented in your chosen domain in the previous section. In this section, we will be looking at a few key verticals – that is, **fintech**, **healthcare**, **marketing, manufacturing, education,** and **cybersecurity** – that have seen increased AI adoption in order to best demonstrate trends within major areas of AI development through these examples. Getting an understanding of how and why AI was adopted in these verticals can give us promising insights into how AI can be applied in other domains as well. Let's explore the adoption of AI in these key verticals.

Fintech

Perhaps the swiftest and most substantial AI transformation has been in the fintech space, and it's not surprising to see why. AI applied to specific use cases can bring about significant revenues saved or generated when done right. According to a recent report by UnivDatos Market Insights, "*The AI in Fintech Market is expected to grow at a steady rate of around 30% owing to the increasing demand for fraud detection, virtual helpers coupled with easy transactions, and chatbots for quick and instant query solutions.*"

Let's take a look at some of the most compelling use cases of AI for the fintech space that have contributed to its quick adoption.

Chatbots and virtual assistants

Chatbots or **conversational AIs** use **natural language processing** (**NLP**) to handle customer service through what's referred to as **sentiment analysis**. Particularly now, as LLMs and APIs that connect to them are becoming more prevalent, these capabilities are becoming even more accessible. These language models are used to find patterns in what end users are asking for to inform things such as a centralized knowledge base, a frequently asked questions page, or to help the company understand more about their customers and what they're looking for in the form of customer surveys. Where ML and deep learning work with data points, in NLP, the data points are the words themselves and it's their orientation and arrangement that gets optimized.

Players in the fintech space are able to use chatbots to also optimize for other things, such as customer feedback with their fintech apps, process automation, and reducing wait times. They're also great for capturing generational sentiments, such as not wanting to jump on a quick call, something that's increasingly prioritized by the younger generations, who prefer to interact with a brand digitally. **bank of america** (**boa**) has seen massive success with its chatbot, Erica, and plans to increase its AI capabilities because of how accessible Erica has been to its users. Erica can assist users with tasks like managing accounts, tracking spending, and providing personalized financial advice. Launched in 2018, Erica has surpassed 1 billion interactions and helped BoA spike earnings by 19% since the launch. Over 32 million clients have benefitted from Erica's insights and notifications. Other examples include Cleo, which uses chatbots to help customers analyze spending habits and provide financial insights, and Sage's Accounting Chatbot, which helps customers using its accounting software with financial queries and tasks.

Though the upsides can be significant, chatbots can have their limitations. Potential miscommunications due to language nuances can cause frustration with customers. Though chatbots' 24/7 availability is certainly helpful in a pinch, they're reliant on predefined scripts, which makes complex queries hard for them to service. This has given rise to AI-powered virtual assistants, which are essentially chatbots powered by generative AI models that work in conjunction with NLP models.

Fraud detection

This is a big one, particularly because so many financial institutions are left with a risk when there is fraudulent activity on an account. Therefore, it makes sense that one of the primary use cases for AI that caught on heavily within fintech has to do with catching fraud more efficiently. In this use case, AI is used for anomaly detection within accounts so that banks and other financial institutions can be alerted faster when activity that's outside an existing customer's pattern of behavior is identified. This is also used for money laundering and other illegal activities.

Fraud detection as a use case isn't quite as straightforward as building a chatbot because there are likely multiple layers of AI tools working together to make it happen. It's more likely to be a combination of things:

- First, some version of continuous data mining is at play to discover whether there are overarching, easy-to-detect patterns in the customer data.
- There's also likely to be a rules-based system that's used to mine as it's scouring the data and looking for anomalies.
- On top of this, there are likely unsupervised ML models used to look for patterns beyond the mining and to group activities into clusters to be analyzed later.
- There are often neural networks used as well to learn from established suspicious patterns that do turn out to be fraudulent.

Because of the level of sophistication that's out there now in the cybersecurity space, fraud can come from multiple sources. Fraud can come from a customer themselves, someone impersonating them, an adversarial bot attack, phishing attacks, and other scams. Fintech institutions need to find a way to address all these potential cybercrimes through the help of various layers of AI. Setting up this kind of in-depth operation is expensive, but the cost to run a system like that is nothing in the face of the $51M fintech loses to fraud every year. With regard to fraud detection, there are also temporal qualities to it, so fintech often looks at finding fraud that has occurred, finding fraud as it's occurring, and finding and stopping fraud before it happens.

Companies like PayPal use AI for fraud detection by leveraging ML, DL, and graph technologies to analyze transaction patterns and detect anomalies that could indicate fraud. Mastercard uses predictive analytics, ML, and a proprietary recurrent neural network in their decision intelligence capabilities to assess risk levels of transactions in real time. Zest AI uses ML and advanced analytics to enhance their credit risk and fraud detection capabilities for lending.

Though companies using AI for fraud detection are able to save on potential losses, there are challenges with these applications as well. Balancing false positives with false negatives is a constant battle. Favoring false positives could mean leaving customers frustrated when they make real purchases and favoring false negatives means they would have to spend more money on recuperating costs. Also, as fraudsters get smarter with new, evolving fraud tactics, businesses need to constantly adapt to match their evolved skills.

Algorithmic trading and predictive analytics

AI is helpful when it comes to the fast trading speeds and improving accuracy that's needed to compete with things such as **high-frequency trading** (**HFT**). ML models are used to effectively predict market movements more quickly to help algorithms make their bids, anticipate the best times in the day to make trades, and use historical data for prediction models to understand when the prices might go up. Apart from the main motivator being to make better, more fortuitous trades for their customers, another big driver of the adoption of this AI use case is that algorithmic trading also helps limit mistakes coming from emotional or psychologically stressed traders to limit trade volatility.

Though algorithmic trading relies on a set of instructions to execute decisions, and there are a number of solutions out there that do just that, the underlying technology that supports it is based on predictive analytics. Fintech players are able to crunch so much data and they're so committed to improving their ML models that even small percentages of improvement in their accuracy or precision could result in millions of dollars of saved revenue or maximized profits. Whether they're a start-up that's looking to convert a lead or a bank that's looking to offer a loan based on a credit score, fintech is using predictive analytics to constantly learn from new and historical data to power most of their decisions so that they know how much a potential customer transaction is worth to them in the long run.

Companies like Goldman Sachs use ML and DL to analyze vast amounts of market data to identify trading opportunities, find optimizations in trading strategies, and predict stock price movements. AQR Capital Management uses machine learning and quantitative analysis to optimize portfolio construction, manage its investment risks, and assess the value of various investment factors. Citadel Securities recently partnered with Google Cloud to optimize their quantitative research technology platform, which uses ML, advanced analytics, and time series analysis to inform real-time trading decisions and manage risks as well.

Though leveraging AI for algorithmic trading and prediction is beneficial, managing data quality standards from disparate data sources is a constant challenge. All models have their limitations and this is especially true with unpredictable financial markets that are often responsive to unforeseen events and the chaos of the real world. The speed of market changes requires model adaptation that's frequent and rapid to adjust to a changing landscape. Financial institutions leveraging AI/ML in this way need to remain competitive and innovate constantly to keep an edge in this highly crowded space.

Healthcare

We're constantly hearing about the staffing shortages that plague most hospital and health networks, and healthcare is another reasonable vertical that's embraced the capabilities and help of AI with open arms. Staffing shortages aside, AI has been able to help with one of the most costly impediments to perfection: human error. Because of the increased demand for personnel and the life-and-death urgency that comes with a lot of the medical activities that go on in the name of our well-being and health, AI has been a great contender for fighting some of the most important battles in the healthcare space. We've seen AI applied in a range of beneficial ways, from image recognition and diagnostics to drug discovery.

Imaging and diagnosis

Supervised and **unsupervised** ML models are used to identify, group, and orient medical images to better understand, for instance, which cells in an image might be cancerous and dangerous. Also, because of the digitizing of medical records, there's a large amount of data that needs to be analyzed and learned from. Mountains of data in the form of images from MRIs, CT scans, cardiograms, and ultrasounds are all available to be analyzed en masse. AI is able to help make sense of these images that exist, as well as learn from new graphics and images that are produced as people continue with their medical activities. The more examples it gets, the better it gets at detecting anomalies. As time goes on, we should see improvements in models optimizing further, to the point where most diagnoses will come from AI-assisted doctors.

Because so much data is being analyzed and the use cases for these models are so specialized, the patterns and relationships that the algorithms are looking for are getting refined for accuracy and speed over time. This is doubly true when we consider the rise of diagnostic and treatment apps themselves. Image processing is getting really good at picking up even small blemishes and abnormalities that human eyes can't pick up, so as doctors and patients continue to use apps to assist them on their journey together, and as that data is collected and centralized further and further, we're likely to see more accurate diagnoses over time as well.

We can see practical examples of this in IBM Watson for Oncology, which uses ML and DL to analyze medical images, assist with diagnostics, and recommend potential treatment options for patients. Companies like Tempus use DL to analyze CT scans and MRIs to find and characterize tumors, as well as to assess their size and spread. They also assist healthcare providers with insights from imaging and genomic data, which helps inform personalized treatment plans for the unique characteristics of a patient's cancer.

The integration of AI capabilities with legacy medical systems can make these use cases especially challenging because they take a lot of time, resources, and money to execute properly. Also, strict regulatory standards like FDA approvals and HIPAA regulations can create challenges with the data used to train models for these use cases. Then there's the issue of trust, as DL is used in many of the solutions in the market for these use cases, causing a challenge, with even healthcare providers suspicious of how these models come to a decision. Often, there can also be delays in getting AI models approved through clinical trials, which poses a challenge for the DL models used, which need to test on real-world clinical data.

Drug discovery and research

Drug discovery and research have historically been notoriously time- and energy-consuming. A process that might have taken years or decades in the past can be sped up today with the help of AI. We all collectively saw AI applied in this way to address the COVID-19 pandemic and speed up vaccine discovery, and according to the National Library of Medicine, *"AI is being successfully used in the identification of disease clusters, monitoring of cases, prediction of the future outbreaks, mortality risk, diagnosis of COVID-19, disease management by resource allocation, facilitating training, record maintenance and pattern recognition for studying the disease trend."*

Companies like AbCellera use ML to predict binding affinity and to develop monoclonal antibodies in drug discovery. They analyze data from immune responses to identify potential antibody candidates based on sequences and structural information that can bind to proteins associated with various diseases. BenevolentAI uses ML to repurpose drugs by using historical data and biological insights to find new uses for existing drugs to help speed up the process of developing new solutions. They also use ML to identify novel drug targets by analyzing relationships between genres, proteins, and diseases, as well as to optimize clinical trials by identifying suitable patient populations. BioXcel Therapeutics is thriving in the drug discovery space by using their BioXcel AI platform to analyze genomic, proteomic, chemical, and clinical data to identify potential drug candidates by predicting their safety profiles and efficacy. They use DL to mine scientific literature, clinical trials, and biological databases to prioritize targets for future investigation for solutions to central nervous system disorders as well. They also use computer vision to analyze images of high-throughput screening to identify cellular responses and evaluate the effects of compounds on cell morphology for drug candidates.

Challenges of AI applications in the drug discovery and research space relate to regulatory compliance issues since the regulatory landscape in AI drug development is still evolving; compliance can be difficult to achieve and regulatory standards are complex and time-consuming. High computational needs also make it hard for smaller organizations to play a role in AI-assisted drug discovery and research. Similarly to algorithmic trading AI applications, a changing biological landscape means that models may become quickly outdated. Biological factors of diseases are complex and can evolve over time, making it hard for AI/ML models to adapt.

Marketing – segmentation

Companies have long been using data to capture our minds and hearts through marketing efforts, but today this is further amplified with the use of AI. Targeting the right audience with the message they most want to receive at the right time is an art form and a lot of data is involved with refining that art. Because marketers are often trying to establish a baseline of behavior for conversions, and how they can improve on that behavior, this is an area that's ripe for AI implementation. While marketing use cases see AI applied in a variety of ways, from targeting, advertising, sentiment analysis, content generation, and analytics, we will focus on **segmentation** specifically as it's an area that underlies almost every facet of marketing.

Segmentation allows not just for grouping customers you already have, but a grouping of your prospective customers as well. Segmentation is used for targeting advertising by analyzing large consumer datasets to establish preferences and patterns across consumer groups.

AI can help establish these groups, gather insights and preferences from these groups, and even help generate content relevant to these groups as well. All data sources can be relevant here, including demographics, browsing history and cookies, purchase behaviors, and interactions from social media engagement, and can all be factored in to create various profiles of potential users/customers based on this data.

This kind of segmentation also allows personalization to happen at scale. If you can use profiling data to understand your users, it follows that you can then help bolster engagement with your brand by taking segmentation a step further into personalization. Generative AI models and NLP now make the generation of personalized content incredibly fast. They can even help inform the marketing strategy overall by tweaking language around product descriptions, email campaigns, social media content, and blog content as well.

Companies like Sephora use clustering algorithms and regression models to segment customers based on purchasing behavior, preferences, and engagement levels to deliver targeted marketing messages to specific groups. They also use neural collaborative filtering for personalized recommendations. Spotify uses content-based collaborative filtering to understand users' music preferences and suggest new music recommendations and personalized playlists for segmented users. H&M uses clustering algorithms and predictive modeling to analyze customer data with trend forecasts to align their marketing strategies with their inventory levels for various segmented groups.

Challenges with AI/ML applications in marketing and segmentation include ensuring compliance with data protection regulations like GDPR, CCPA, and others, which can have legal implications but also negatively impact consumer trust. Data quality issues can also lead to marketing campaigns that are irrelevant, which can lower engagement rates and negatively impact the brand. Also, in many cases, segmented users may find themselves in an echo chamber of sorts, with their AI recommendations only showing them certain categories of content, which can limit their experience of discovering new preferences. Personal preferences are always subject to change or be impacted by pop culture and other outside influences, so marketing teams that employ AI applications need to find a way to capture new evolving interests even when segmentation goes well.

Manufacturing – predictive management

We see the pragmatic and utilitarian approach of AI in manufacturing contexts. Here, AI is being used for everything from quality control to supply chain to robotics and the end goal is to optimize the use of time and resources. Particularly today, when global supply chains can be sensitive to fluctuations in the market and global geo-political influences, AI can be used to smooth disruptions.

One of the advantages of the manufacturing industry is it's able to generate a lot of its own clean data to learn historical trends from when the infrastructure is set up properly. AI can learn from real-time production data to understand where there are inefficiencies and bottlenecks so that management can address them. Based on inefficiencies AI systems learn from, they can then help predict when there may be a potential breakdown somewhere along the process. Because so much of the equipment used today might have sensors or IoT devices, we can also predict when certain parts of the factory line might be in need of maintenance as well.

Predictive maintenance also helps advance the overall management of the facilities themselves. Managers can use the predictive capabilities of AI to better plan how they allocate resources, plan for downtimes and disruptions to their assembly lines, budget for maintenance costs and investments, and track the overall productivity of their operations. This doesn't just bolster productivity gains but also give peace of mind from being proactive and managing things in a planned, controlled way.

Companies like Avathon specialize in AI-powered solutions for predictive maintenance, optimizing manufacturing processes, and detecting early signs of equipment failures. Their suite of AI products for manufacturing, known as the Industrial AI Suite, uses AI/ML to predict potential machinery failures before they occur, helping to minimize unplanned downtime and maintenance costs for the oil and gas, aerospace, and manufacturing industries. Uptake, a predictive maintenance company that specializes in trucking fleets, has built their platform around risk management and downtime minimization as core value propositions of using their product.

Education – personalized learning

We all have different writing styles and a one-size-fits-all approach to learning makes it so that most students end up with sub-optimal learning experiences. Education is ripe for a boost from AI to help learning paths and strategies become more widely accessible to students from many diverse backgrounds. This is encouraging to see, particularly since today's reality of working in education has already alienated many well-meaning educators in the US and across the world. Students are in dire need of learning pathways that meet them where they are and can grow with them, without adding additional strain on their already overworked teachers!

Today, there are so many use cases of AI to help education platforms adapt to the learning styles of the students they are serving. Understanding the preferences, weaknesses, and strengths of the end users/students contributes to crafting personalized learning experiences for individual students. The outcomes of these customized learning paths are that learning can be self-paced, content can be tailored, and gaps in knowledge that need to be improved can be prioritized for each student. In many cases, this also adds an additional premium of shielding students from potential judgment from their peers since students can learn confidently without being singled out. Adaptive learning also makes it so that students can progress according to their ability to receive learning. Difficulty levels can be adjustable, personalized practice exercises can help fill in gaps, and underlying ideas can be continuously reinforced as a student progresses along a path.

Companies like Duolingo use an LSTM (RNN) model for speech recognition and feedback on pronunciation as well as reinforcement learning to tailor lessons based on how well users have performed on previous exercises and lessons. Squirrel AI uses knowledge graphs to establish relationships between concepts and track student mastery, as well as reinforcement DL to predict student learning outcomes and optimize difficulty levels for problems.

AI applications in education can face some unique challenges. If learning paths are too personalized, they may actually dissuade students from learning because they may favor their weak areas and overwhelm them with difficulty. Learning content also has to adhere to local, national, or international localization and cultural standards. Bias also needs to be a priority so that models are not trained on data that underrepresents the learning group, which could have unfair or unequal learning outcomes for some students. Even teachers could pose a challenge, particularly if they are AI skeptics or fear job displacement.

Cybersecurity – anomaly detection and user and entity behavior analytics

Cybercrime is up, and cybersecurity continues to be one of the most prevalent sources of AI/ML adoption that we see most regularly because of the persistence and sophistication of cyber criminals. The nature of cybersecurity is ever-changing, which means that, over time, companies need to keep step with their potential adversaries. It feels like we hear about a new data leak every day, and perhaps the most alarming part of cybersecurity is the number of attacks you *don't* hear about. In this section, we will discuss two popular use cases used in cybersecurity: **anomaly detection** and **user and entity behavior analytics** (**UEBA**).

Uncovering patterns and changes in those patterns is at the very heart of anomaly detection. In cybersecurity, this established baseline and deviation from it are what create the use case of anomaly detection. However, anomaly detection is quite broad in scope because it can refer to identifying outliers or abnormal events in any system, network, or dataset. Areas in cybersecurity can range between network traffic and system logs and even user behavior. Any irregular traffic spikes, unusual data transfers, or system behaviors that don't align with historical data would be good contenders. Anomaly detection is often used in network security to combat DDoS attacks or potential data breaches. Once there's been an anomaly and an action is required from the system, we can then move toward rectifying it somehow. Often, cyber attacks come from within networks and have clever ways of hiding their tracks, but with advanced pattern recognition that's used for anomaly detection, AI systems can see when some actor/user is behaving in ways that don't make sense for a normal user.

As we went through some of the use cases for various verticals, you're likely seeing patterns of your own emerge in terms of some of the underlying tech used to power those applications of AI. One of these overarching use cases is encapsulated in UEBA, which is, in many ways, an undercurrent for many of the use cases we discussed. As an AI PM, it will be really helpful if you can understand the power of UEBA and how to apply it within your product and beyond, particularly as you're in the operational phase of building out increasingly powerful features over time.

You can think of UEBA as a network of insights that are gathered by basically pooling all your users' actions and subsequent data that is generated from those actions on a daily basis. Once all this data is centralized and analyzed, it's able to offer you profound insights by giving you a baseline for not just individual users but for all your users all at once as well. This makes things such as finding anomalies or new patterns, no matter what the change is, easier because you can establish triggers or actions that alert your user or your internal workforce when something that requires attention has happened.

Companies these days need to stay more vigilant than ever, and many don't have the headcount or allocated budget to truly stay on top of this potential risk, which is why the breaches keep coming. This is a big part of the reason cybersecurity has adopted AI so readily: it's ripe for it. When use cases are notoriously underfunded and understaffed, AI has an opportunity to get adopted and truly shine.

Companies like Darktrace use a combination of supervised and unsupervised ML to model the behaviors of every device, user, and network within an enterprise through pattern recognition. Varonis is used by companies in the financial sector to prevent sensitive data leaks. They use ML and behavioral analytics to find anomalies in data access and usage in order to find compromised accounts or insider threats. They also use NLP for customers to investigate and analyze breaches through the Athena AI (generative AI) layer of their data security platform. Exabeam's advanced analytics platform uses machine learning to create detailed timelines of abnormal user behavior for threat detection and access to sensitive data as well.

This kind of behavioral profiling does, however, struggle with false positives as not all anomalies are malicious, and even routine activities like system updates could be flagged. In general, these applications of AI generate too many alerts and alert fatigue is a user pain point that these companies have to grapple with. Too few alerts and you risk exposure, but too many risks disengagement.

Thought leadership – learning from peer groups

At the start of this chapter, we discussed the idea of building a foundation in your domain and understanding (as much as possible) what the specific pain points in your domain are with the aim of building an AI product that will serve the space for years to come. Building a product that works well, aligns with what your customers need, and employs modern technology is a fast way to build the credibility that's necessary to spread thought leadership across your domain.

You might choose to take on a leadership role in your domain in order to be positioned as a leader in your industry. With all the knowledge hubs, white papers, and product one-pagers that are out there, it's common to see companies adopt the role of an industry thought leader for a number of reasons. Maybe this is for marketing, inbound leads, glory, or simply because they're generous with their knowledge and success. Your choice of whether or not to be open with your online communities about your product and proprietary information might even come down to your product and brand strategy. Does it open up pathways for your end users, customers, and fans to follow along with your journey? Could it contribute to a stickier, more loyal customer base?

Choosing how open and transparent you want to be with your company's choices and major roadblocks will be up to you and your leadership team. Will you choose to be an example for the other players in your field? Or, will you choose to keep quiet and hold your algorithmic secret sauces, dirty laundry, and success stories close to your chest? Being a true thought leader in an industry may come at a financial and design cost if you're not careful of how much information you let out. As with all risky endeavors, sometimes we need to open ourselves up to risk in order to also accept the gifts that come with transparency.

As we've mentioned so many times, this period of applied AI products is still new. Even with all the buzz coming from the generative AI space, we're still just scratching the surface of how companies and users will adopt these technologies. In the end, we don't know how useful and groundbreaking they will be until we see them prove their worth in action. We have yet to understand the future use cases that will allow companies the greatest growth and success. The spirit of innovation and technological advancement is fed through example. Every use case and applied example has the ability to inspire us to discover a new use case. The newness of AI brings with it a kind of excitement about what's possible.

Particularly in a time when so many are apprehensive about whether AI can and should solve all our biggest problems, maybe it's up to companies, PMs, technologists, and entrepreneurs to embody the collaborative quality of this spirit. AI products bring together science, data, technology, and humanity in a way that has shown so much promise already. Should you choose to open up and share the examples that do come up in your experience, whether directly or indirectly, the gifts that come along with that won't just benefit you or your organization; they will benefit us all and, foremost, serve the highest ideals of your business: to best serve your customers and market.

Case Study

The market

Given the size and resource constraints of Waterbear as a company, our case study would firmly be in the **Visionaries** category. Here, Waterbear's greatest strength is that it's an app with a lofty goal: to deliver robust mental health support to its users not through outside sources but through internal reflection. Though it's an app you can freely share with a therapist, a friend, a partner, or any other external third party, it is meant to be a personal journey through your own journal prompts. Waterbear's goals were quite visionary from the onset because its leadership team didn't see NLP being leveraged in the form of a personal dashboard in the user experience for any of the other competing products out there. This sets it apart from its competition, which typically serves as either a simple mood tracker or as a connector to mental health networks and therapists. Bringing a holistic experience to its users meant that they could track their progress from week to week, month to month, or year to year. The sustainability of long-term advancement toward health, goal setting, and accomplishment is something that the founders of Waterbear felt was sorely lacking among the mental health and wellness apps that were out in the market. Depth was also lacking. Building an app experience that could go beyond capturing a mood but could capture a wide range of human-centric goals and pursuits wasn't easy to build, but it was worthwhile enough.

Product design and strategy

For Waterbear, the *Why AI* question was an easy one to solve. You need AI that's good at working with unstructured, unlabeled textual data. Though some aspects of the product experience were able to be established with the end user, the most substantial work the NLP algorithms were running in the background was around the establishment of goals, sentiment analysis, and tracking progress.

It takes AI for that kind of work because every person that's using the app is different. They say things differently, words mean different things to different people, and their reactions or emotions about certain words will even vary from day to day. Training an AI that can capture that kind of elasticity, and will be able to notice drift and fluctuations from day to day is important. With their product Akeira, Waterbear had to zero in on nuances in language that pertain to specific demographics of women. In this case, the demographics weren't just specific to ages or racial backgrounds but even generational identities and subcultures.

For this reason, they focused on training each model for each user and using the combined data for overarching trends and insights and not to inform model performance. This meant that they were customizing their own proprietary pre-trained language models and cloning those models for further training on the input data coming from each end user independently. Because the journal prompt inputs were so idiosyncratic, keeping the training process separate and optimized at the per-user level meant that they could capture that nuance in the personalized dashboard.

Thought leadership

Since transparency, trust, and authenticity are important to Waterbear, they were authenticity-forward with their product strategy. The Akeira product team transparently talked about how they built and trained their models and they shared that information with their user groups and online communities. They focused on a grassroots marketing campaign to maximize organic users and found that the more transparent they were able to be about their own challenges, the more their customers and end users trusted them.

Akeira's AI PM maintained a product blog on the company page where they frequently discussed important product decisions they felt comfortable sharing with their online community to better represent their choices and decisions to a wider audience. This fulfilled a few key objectives: it widened their circle of feedback, brought awareness and discussion to difficult and sometimes controversial decisions, generated brand awareness for Akeira and Waterbear overall and, most importantly, positioned the company as a thought leader. Their AI PM spoke at conferences and industry events often and this broader visibility of Akeira garnered the attention of industry awards and recognition. All of these activities generated significant amounts of inbound leads and inquiries about the work they were doing.

Summary

This chapter was geared toward markets, positioning, and common use cases of AI/ML products. We've been able to look at how AI can be optimized for certain domains and markets and how AI can commonly be leveraged in various verticals that are seeing a high saturation of AI products. Through those use cases, we've been able to see how companies leverage AI to be able to make the most of the data they have. As an AI/ML PM, you won't be building your AI-native product in a vacuum. You'll regularly be studying your market and your competition to make sure you're bringing use cases for AI that truly set you apart.

In the next chapter, we will discuss design principles and elements that help set up an AI-native product for success.

References

- *AI and the COVID-19 vaccine: Moderna's Dave Johnson*: `https://sloanreview.mit.edu/audio/ai-and-the-covid-19-vaccine-modernas-dave-johnson/`
- *The role of artificial intelligence in tackling COVID-19*: `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7692869/`
- *FinTech execs expose the real costs of Fraud*: `https://www.pymnts.com/money-mobility/2022/fintech-execs-expose-the-real-costs-of-fraud/`
- *One million people are now using Erica - BofA's AI-powered chatbot*: `https://netfinance.wbresearch.com/blog/bank-of-america-ai-powered-chatbot-strategy`
- AI in fintech market predicted to reach $49.43 billion by 2028 with evolving fraud detection technologies: Industry analysis through 2023-2028: `https://www.globenewswire.com/en/news-release/2024/02/08/2826134/28124/en/AI-in-Fintech-Market-Predicted-to-Reach-49-43-Billion-by-2028-with-Evolving-Fraud-Detection-Technologies-Industry-Analysis-Through-2023-2028.html`
- *Gartner's Magic Quadrant*: `https://www.gartner.co.uk/en/methodologies/magic-quadrants-research`
- *From chatbots to visionary generative AI-powered productivity assistants*: `https://www.sage.com/en-gb/blog/chatbots-ai-automate-powerful-results/`
- *Cleo*: `https://web.meetcleo.com/`
- *A guide to leveraging data analytics in fraud management*: `https://www.paypal.com/us/brc/article/data-analytics-fraud-management`
- *Artificial intelligence at PayPal – Two unique use-cases*: `https://emerj.com/ai-sector-overviews/artificial-intelligence-at-paypal/`
- *Discover Mastercard's Security By Design resources tailored around detection topics*: `https://www.mastercard.com/globalrisk/en/resources/all-resources/detect.html`
- *Mastercard jumps into generative AI race with model it says can boost fraud detection by up to 300%*: `https://www.cnbc.com/2024/02/01/mastercard-launches-gpt-like-ai-model-to-help-banks-detect-fraud.html`
- *Proven AI for a thriving lending ecosystem*: `https://www.zest.ai/`
- *Leveraging AI in electronic trading: A game changer for investment banks*: `https://electronictradinghub.com/leveraging-ai-in-electronic-trading-a-game-changer-for-investment-banks/`
- *AQR's problem with machine learning: Cats morph into dogs*: `https://www.institutionalinvestor.com/article/2bsw8i4p5x9et9ny1mc5c/corner-office/aqrs-problem-with-machine-learning-cats-morph-into-dogs`

- *How Citadel Securities is reimagining quantitative research on the cloud*: https://cloud.google.com/transform/citadel-securities-reimagine-quantitative-reseach-cloud-scale-speed
- *IBM Watson for Oncology*: https://www.ibm.com/docs/en/announcements/watson-oncology?region=CAN
- *AI-enabled precision medicine*: https://www.tempus.com/
- *AbCellera and Lilly slash antibody selection time for COVID-19 with AI/machine learning*: https://www.biospace.com/abcellera-and-lilly-slash-antibody-selection-time-for-covid-19-with-ai-machine-learning
- *AI-enabled drug discovery*: https://www.benevolent.com/
- *AI-based drug re-innovation*: https://www.bioxceltherapeutics.com/ai-based-drug-re-innovation/
- *Personalized recommendation on Sephora using neural collaborative filtering*: https://himaeda.medium.com/personalized-recommendation-on-sephora-using-neural-collaborative-filtering-1d4c1fa3f98d
- *Inside Spotify's recommender system: A complete guide to Spotify recommendation algorithms*: https://www.music-tomorrow.com/blog/how-spotify-recommendation-system-works-a-complete-guide-2022
- *Customer segmentation analysis for H&M*: https://dwang2017.wordpress.com/wp-content/uploads/2017/01/hm-customer-segmentation_report_team2.pdf
- *How Duolingo's AI learns what you need to learn*: https://spectrum.ieee.org/duolingo
- *China has started a grand experiment in AI education*: https://www.technologyreview.com/2019/08/02/131198/china-squirrel-has-started-a-grand-experiment-in-ai-education-it-could-reshape-how-the/
- *Darktrace AI: Combining unsupervised and supervised machine learning*: https://cdn.prod.website-files.com/626ff4d25aca2edf4325ff97/62a292c20b7371fb41311930_wp-machine-learning.pdf
- *Data Security Platform*: https://www.varonis.com/products/data-security-platform
- *Varonis introduces Athena AI to transform data security and incident response*: https://www.varonis.com/blog/varonis-introduces-athena-ai
- *AI SIEM: How SIEM with AI/ML is revolutionizing the SOC*: https://www.exabeam.com/explainers/siem/ai-siem-how-siem-with-ai-ml-is-revolutionizing-the-soc/
- *Avathon*: https://avathon.com/
- *Keep your trucks on the road*: https://www.uptake.com/fleet

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

9

Product Design for the AI-Native Product

So far, in Part 2 of this book, we've covered a lot of areas as they pertain to building and maintaining an AI-native product. But we still haven't covered **product design** and the extent to which PMs should be well versed in design. In this chapter, we aim to do just that: understand the design elements that best set up a native AI product for success. We will go into design principles and the priority levels that are important to building any product, and we will home in on which of these elements will be important to focus on when you're setting out to build an AI product natively from the ground up.

Often, when we discuss the rise of AI, we imagine a world where a number of human tasks will be automated. This has left a lot of people wondering where our place as humans will fall among all the technological advancement that's been going on, particularly in recent years. Not all product designers will know how to be product managers, but all product managers will have to have some understanding of design if they're going to do their jobs effectively. Design is, perhaps, the most exciting and glamorous part of the discipline of product management, where function, creativity, and user experience converge. Solving complex, human-centered problems isn't something that's easy to automate, and your understanding of design will have a huge impact on your product.

Designing a product isn't just about being creative. It's also about how you can craft an experience that not only makes the lives of your end users easier or more joyful but also communicates the value your product can give others in a meaningful way. Product design is psychological, interdisciplinary, and all-encompassing. It's not just about making a product that works well from a utility standpoint; it's also about building something that emotionally resonates with your end users and has them coming back for more. Throughout this chapter, we will:

- Go over general product design foundations
- Examine how those foundations are challenged when we throw AI into the mix
- Learn how to use tools that help us balance design priorities
- See how we can weave those priorities into the storytelling that accompanies our product

We will finish with a product design overview of our case study example as we've done with the other chapters in this section.

Recap of common product design terms

Customer journey: The complete series of interactions a customer has with a company or product, from initial awareness to product onboarding, mapping key thoughts, emotions, pain points, and touch points until purchase.

User journey: The complete series of interactions a user experiences with a product, including discovery, engagement, and ongoing use, mapping key thoughts, emotions, pain points, and touch points at each stage of the life cycle.

User interface (UI): The UI encompasses all visual elements and interactive features (screens, buttons, icons, menus) of a digital product, focusing on aesthetics, layout, and visual elements.

User experience (UX): The UX encompasses the entire user experience, including UI, functionality, accessibility, and overall ease of use, satisfaction, and emotional responses.

UI vs. UX design: While UI design focuses on the visual representation and surface-level interactions of a product, UX design encompasses the entire user experience, including look and feel, user research, content structure, and interaction flow, to ensure a seamless and satisfying experience.

UX research vs. market research: UX research is product-centric, focusing on user behavior, needs, and motivations to improve product interaction. Market research is broader in scope, focusing on the market landscape, competitors, trends, demographics, and preferences to inform overall business strategy and product positioning.

Customer journey vs. user journey: While customer journey encompasses the broader pre- and post-purchase experience, including marketing, customer service, and brand perception, user journey focuses on the specific functional and task-oriented experience of navigating and using a product.

Product design elements 101

In *Chapter 6*, we covered the stages of AI product development. You'll likely see concepts mirrored here, but it's important to note that developing a product and designing a product are two different things. Certainly, there's an interplay between the two; both are required to bring a product to market. Product development refers to the entire process of building a product. Product design focuses on the creation of the product's user experience and form itself. If you refer back to *Chapter 6*, you can think of product design as one of the areas that will be covered in *Phase 1* (ideation) and *Phase 3* (research and development) of the product development process. As you iterate on your ideas and reinforce them with research and feedback, you're doing the work of maturing your product design.

The following is a visual guide to organize your product design elements and refer back to. Keep in mind that these stages may change as you're building, but they generally take the following sequence. As you're navigating the task of bringing AI-native products to market, you'll revisit various stages at various times.

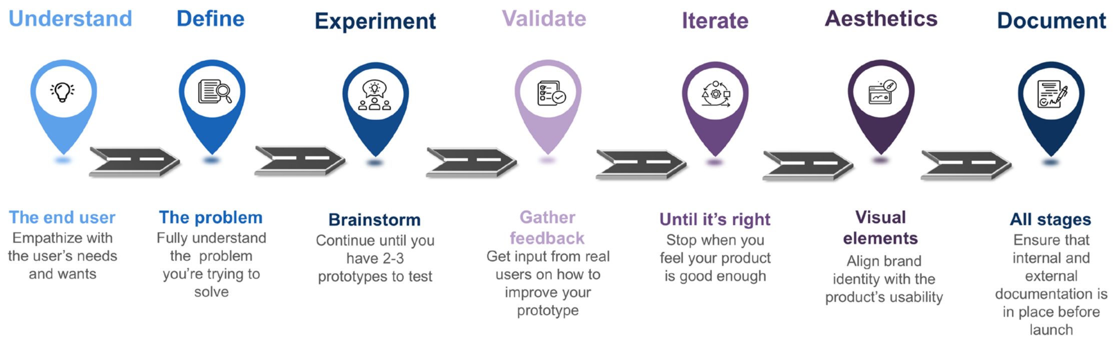

Figure 9.1: Product design elements

🔍**Quick tip**: Need to see a high-resolution version of this image? Open this book in the next-gen Packt Reader or view it in the PDF/ePub copy.

🔒**The next-gen Packt Reader** and a **free PDF/ePub copy** of this book are included with your purchase. Unlock them by scanning the QR code below or visiting `https://www.packtpub.com/unlock/9781835882849`.

In the upcoming sections, we will explore these facets of product design and how they relate to each other further. Because this chapter is focused on AI products being built natively for the market, you'll likely have to go back and forth between stages until you get it right. You can treat each section as a milestone to bring your product and leadership team together because succession from one stage to another is a milestone worth celebrating! All these steps should be done as collaboratively and cross-functionally as possible. Make sure you bring a diversity of voices into the room when you do.

Understanding the end user

There's a reason the product design process starts with the user. You need to go through the journey of getting to know what your users actually need and want. They may not even know what that is themselves. You might have some idea of what they already want and need without even talking to them. That idea might be wrong and you need to set your ego aside enough to believe that you may, indeed, be wrong. So, it really is a journey you're going on where you need to deeply empathize with the person you're building for. Start with the things you know for sure.

It's hard to build anything for anybody if you can't understand what's most important to them. Lots of people want to build a product that everyone loves, but in order to bring a product to market in a way that resonates with a lot of people at once, you need to understand what drives and annoys your customer the most. That means a significant amount of work needs to go into understanding your end user spectacularly well. Starting with understanding your end user early on, before your idea of the problem they're facing sets in, helps avoid bias and assumptions that could be costly to integrate downstream.

What this looks like is user research, potential customer interviews, and market research to understand the target audience your product will serve. Even if you think you understand this end user and you're already mentally crafting an incredible potential solution for their problems, you still need to put in the necessary work to validate your understanding. Take the time to understand their behavioral and usage patterns, as well as their expectations and preferences. No matter how sure you may be, without actually talking to your potential end users, you may not be building the solution they need but the solution you want to build. Use this period as an opportunity to get to know their response and feedback to existing solutions on the market to understand how they are falling short as well.

Some questions you can ask potential end users at this stage are:

- What is the biggest challenge you face in your current workflow or daily routine?
- What are your existing options for solutions to address those challenges?
- What do you like and dislike about those existing options?
- Is there a "perfect solution" and, if so, what would that look like for you?
- How do you typically discover and decide on new products?
- What are your core user needs?

Defining the problem

Once you've amassed enough information about your customers and end users, you'll start to group these findings and establish the actual problem you're trying to solve. In many cases, these first two steps are inverted when the end user's perspective is brought in after the problem has been considered. Technologists, product managers, and well-meaning entrepreneurs can jump the gun on this one pretty quickly. They might already be veterans in the space they're trying to bring a new product to market in. They may already feel they understand the problem and then validate that problem with evidence from the end users they interview, if they do this at all. In some cases, they may make an assumption about the problem they're solving and cherry-pick user interviews that only support their underlying assumption.

Doing this might satisfy your ego temporarily, but it will create more work for you downstream when you discover the problem is actually more complex or multifaceted than you initially thought. The work of designing a product balances art and science. Sure, you're creatively thinking about how a product needs to come together to suit the needs of your customers, but you also have to gather evidence and see how that evidence stacks up against your current understanding of the problem. Feel free to let yourself be guided by the results of your user research. Leave enough room to be pleasantly surprised by the findings you discover. You don't need to know everything. But you do need to make sure you respect the guidelines of product design and build with end users in mind.

The following are some key questions that you should be able to answer:

- Can you define the specific problem in three sentences? Can you define it in one?
- How do you know this problem is significant and worth solving?
- Does the end user analysis support the weight of this problem?
- What are the different facets of this problem? How are those facets experienced by various end user segments?
- What assumptions have been made about the problem and how have they been validated or disproven?

Experimentation

Now comes the fun part. Experimenting with ideas and concepts of what to include in your product allows you to take the data you've gathered so far to understand your end users and validate your problem and turn it into a form you can work with and mold. You won't want to rush this step. This is where you can create a series of brainstorming sessions and workshops that will inform your eventual prototypes. Let your creativity run free here and make sure the team you're building with feels comfortable enough to bring up ideas that aren't very good. In fact, don't even worry about whether those ideas are good or not at the beginning. Let your team's creativity run free. The only restraint to include here is to make sure that all the ideas actually address the customer needs identified in the first section. For more information on team roles and responsibilities, feel free to refer back to *Chapter 6*.

You can discard ideas that don't work so well later on when you've landed on a few options you feel confident enough in to move on to the prototyping stage. This stage is about excavating the corners of the mind for all the ideas and concepts that your product stakeholders and designers feel compelled to consider when building the product. Let their structure be loose because this is the stage of product design that's the most free-flowing. You want everyone involved to be in the flow state, which means you want them to be thinking visually and experientially. Think here in terms of whiteboard exercises, loose mockups, sketches, and mood boards. You want to get all the ideas that occur to your team out and make sure they're expressed first. Then you want to bring everyone together to think more critically about these ideas and agree to 2-3 prototypes at least.

Once you have a few product ideas that are most compelling, you can turn them into the prototypes you'll validate and test. Prototyping will take the elements that were included in your brainstorming sessions and thread them together into a few product experiences that you'd feel good about bringing to potential customers. This allows you to make a visualization of your product: a jump from the formless world of ideas into the defined world of form. How formal these prototypes are will depend on your own preferences, but make sure these prototypes capture the visual representation of your product in sketch or wireframe form. At this stage, you haven't actually built and developed anything yet. You're still in the early stages of brainstorming potential solutions to the problem you've defined. It's just that those potential solutions are shareable and you can start to get feedback from potential end users.

You should be able to answer the following questions at this stage:

- What are all the potential ideas and concepts that have emerged during brainstorming sessions?
- How do the 2-3 prototypes you've settled on relate to the customer needs identified in the first stage?

- Are the ideas explored wide enough? Have a wide range of unconventional or "out-of-the-box" ideas and potential solutions been considered?
- Are there ideas that stray too far from the core user needs?
- Do all team members involved in the brainstorming process feel comfortable enough to express their thoughts fully?
- Have all ideas been captured and documented during these brainstorming sessions?
- Are the prototypes you've settled on detailed enough to convey the main functionality but flexible enough to iterate on?
- Do the ideas and prototypes align with the overall product vision and strategic objectives of the company?

Validation

Now that you have something you can put in front of people, you can start amassing feedback that will directly impact the direction and presentation of your product. In this phase, you're putting your prototypes in front of real human users who may one day purchase your product, or represent groups of users who will. This is where you're checking that the feedback aligns with the initial feedback you gathered in the first phase. Here, you're also checking to make sure your prototypes are accurately being perceived, that their usability is being understood even in prototype form, and that your users feel it will solve their pain points. This stage should also give you a few more ideas to help make your product better.

In this phase, you're not only looking to see that your prototype is fitting the bill but also that it's capturing additional complexity you might not have thought of or been aware of in the first round of user feedback. This early in the game, there's plenty of room for improvement, and you'll get an intuitive sense of how many rounds of testing and validation to do. As you continue to mature and refine your product, your testing and validation needs will mature. In some cases, you'll be testing not just the ideas and prototypes, but eventually reliability, performance, and compliance as well. Remember, these phases are design principles, not exactly a set list of sequential steps. You'll be coming back to test and validate with each new iteration that emerges from it.

The key questions you can ask in this stage are:

- Does the feedback from users at this stage align with the initial user research findings from the first phase? Are there significant discrepancies between them?
- Are users understanding the purpose and functionality of the prototype(s) as intended?
- Are there challenges in how users navigate through the prototype(s) to complete core tasks?
- Do users feel the prototype addresses their pain points effectively and efficiently?
- Which features do users find most meaningful, useful, or appealing? Which features do users find most confusing, unnecessary, or counterproductive?
- Have any new ideas, needs, or suggestions emerged from this second round of user feedback during the validation stage? Are there any edge cases, scenarios, or unexpected complexities?

Iteration

Each iteration of your product that comes out of your testing and validation phases will represent a level of your product's maturity and growth. The idea with the iteration phase is to ensure your product is growing in the right direction, that it still resonates with potential end users and customers, and that it is capturing the feedback they're bringing up. The more the real humans who will use your product are in the mix, the closer you will come to designing a product that's intended for them and puts their concerns first. Each iteration will get you closer to a user experience that aligns best with what the majority of your customers and end users want from your product.

Building a user experience that is as seamless and frictionless as possible is the main objective of each phase of iteration you undergo. Make sure your product offers an experience that's intuitive, easy to use, and accessible for your customers. Every time you make new iterations of your product, you should be moving in a direction that makes your end users feel the product is becoming more user-friendly and more emotionally satisfying than the previous versions. Good is the enemy of great. It's hard to know when to stop iterating and arrive at a product design that's good enough. Is "good enough" enough for you and your team? At what point will your final draft become great? When you start to feel your product is meeting your expectations and your customer needs and delighting them, you're getting close. Each iteration you land on should be a sustainable option for production, so this is also where you would make decisions to exclude potential versions of your product that are unsustainable to build.

Here are some key aspects to consider during the iteration stage:

- Communicate the purpose and the limitations of the prototype to users during various rounds of iteration and validation.
- Define what your product looks like when it's "good enough" and when it's "great."
- Decide collaboratively with your team about how many rounds of validation and iteration you anticipate needing before moving to development and what it will take to go to market with this "great" version.
- Define top areas for improvement identified through user feedback and multiple iterations.
- Refine the final prototype to better meet user expectations and needs.
- Set specific criteria about how you will determine when the prototype is ready for development.
- Assess risk based on current iterations for potential performance issues.
- Confirm if the evolving prototype continues to align with the overall product strategy and vision.

Aesthetics

Eventually, you come so close that you're ready to start bringing in the visual elements that polish off the look and feel of your product. This is where the look and feel will start to align with your brand identity. It's not just a product/user experience you're crafting anymore. This is the phase where you're getting your product design market-ready. Aesthetic and marketing choices will need to be made here in a way that aligns with your company's overall aesthetics. At this point, you're pretty sure about the product design you want to go with. You've tested and validated several prototypes, perhaps a few rounds of each, and you've decided which you'll be going to market with. By this point, you've also done your quality testing as well. This is where you'll put in the finishing touches and work on communicating the value and look of your product to the world.

You'll want to align your product aesthetic with the overall brand identity of your product to make sure company values are being communicated as effectively as product functionality is. This phase is about ensuring consistency and cohesion across multiple levels because once your product is out there on the market, perception is everything. Balancing aesthetics with product functionality is a way to communicate to your end users, prospective customers, and the market at large that you've put the work into building a product that works, feels, and looks intentional. Some may feel that aesthetics are secondary to functionality, but even if they are, people will see what your product looks like before they actually use it.

Some questions you can ask at this stage are:

- Does the visual design of the product align with your overall brand identity and values? Are consistent color schemes, typography, and visual elements that are recognizable integrated?
- Do visual elements enhance or distract user experience?
- Are aesthetics and functionality balanced? Does the product look visually appealing and does that make it easier to use?
- Are the design elements used inclusive and accessible for all users?
- Have you identified visual inconsistencies or areas that need to be refined for a more cohesive look?
- How well do visual design elements communicate the value proposition and unique selling points of your product?
- What final visual elements or marketing materials (product videos, landing pages, and imagery) are needed to support the launch of your product?

Documentation

Pre-launch, you'll focus on building enough content to make sure your customers are set up for success as best you can. The work of building this support is ongoing, of course, but you'll want to have enough there to feel comfortable with launching and making sure your company won't be inundated with questions and concerns following the launch as best you can. Documenting your product is important, of course, but it's equally vital to have resources for customers who will be buying your product post-launch. This is where all the content related to your technical documentation, product documentation, marketing, distribution, and customer success knowledge bases will be handled. Remember that this is just a phase of your product design. It's meant to accompany your product and should have a branding, look, and feel that's consistent with your product experience.

This content will also be evolving over time. As you continue to gather more and more customer feedback, particularly post-launch, you'll need to make sure your accompanying documentation and marketing collateral align with your evolving product. Don't think of documentation as something separate from your product because it isn't. Documentation is the mouthpiece of your product and your end users and customers will be relying on it to make sure they understand the capabilities and value of your product. If your documentation and marketing fail to do that adequately, you're not speaking clearly to your potential customer base. Documentation will also be handed off to the development team that's building your product, so make sure you have customer-facing and internal-facing documentation in this phase.

Here are some key questions to consider at this stage:

- What core features and functionalities need to be covered in a customer-facing document(s)? What level of detail are you comfortable including for these documents?
- Have you created clear guides detailing prototype design elements and rationale for development to refer back to when they start development? Is there an internal knowledge base that covers product architecture and validation processes for your development team?
- Do you have a system in place to capture and share prototype upgrades/changes (and subsequent product documentation) after each sprint?
- Are there detailed user stories, workflows, and use cases documented to help relevant teams understand the prototype's functionality and user experience?
- How will documentation issues, gaps, or pain points be handled? How will feedback mechanisms be used to make sure there's a way to receive internal and external feedback on documentation?

What makes the AI-native product design process special?

Given the massive impact AI is having on our workflows and products, we're seeing it have a transformative effect on the design process itself. Here, we're not just seeing AI have an impact on the act of designing products, but on designing AI products intuitively and taking a human-centered approach. What's perhaps most interesting about integrating AI isn't necessarily the tech itself, but the fact that it puts an added premium on the parts of the process that require human intervention. Considering the central role that user needs, interests, and challenges play in product design, this is further amplified when AI is involved in the design process.

In this section, we will cover some of the unique considerations AI products bring to the product design process and how to account for them when designing your AI-native products.

User obsession

Whether it's using AI to better understand and anticipate user behavior, establishing profiles and patterns from past interactions with your product, or discovering new user preferences and priorities, AI products are prime for making quite a lot of progress compared to their traditional software counterparts because of the level of knowledge gathering that's possible with them. In traditional software products, we're still able to gather a lot of data about usage and customer needs. This is the nature of product design, as we covered in the first section of this chapter. But with AI products, we're able to align product design principles with AI algorithms and gain even more insight from the data we gather.

Because the premium of understanding user needs is so high, it's a top requirement for AI products. With more information about users, you're able to design a product experience that's incredibly appealing and "sticky," which means it has staying power with your end users. Amassing this much information about your product doesn't just impact the overall aesthetics, user experience, and form, but it also creates an opportunity where you can create different product experiences for groups of users or even individual users. Your ultimate goal in product design isn't just to create a product that works optimally and looks good; you're also trying to create a product experience that's intuitive and emotionally connected with your end user.

In traditional software products, you might have a threshold where you're comfortable to stop building because you and your go-to-market team might feel the product has reached a level where it's finally fit for general consumption. Sure, you might incrementally add more features and adjust your product design based on ongoing feedback. But by and large, your product experience itself will be uniform. That isn't the case with AI products, where tailored experiences are the norm. In *Chapter 7*, we discussed the idea of "productizing" the ML service. The idea of user obsession in AI product design reinforces that idea because if you have infinite time and resources, and you really wanted to, you could design a different product for each user. AI not only opens up the notion of personalization, but it also stretches the use cases for your product as well.

It's also not just about knowing your users; it's also about interacting with your users. AI products, particularly these days, often incorporate generative capabilities to capture the nuance of their users' needs. This means that instead of having a uniform, predefined user experience and the occasional product survey going out to some customers, you're actually incorporating a lot more opportunities to interact with your end users. Perhaps your product has a chatbot, voice or video component, which we're sure to see more and more products having with the passage of time. You'll be able to get more feedback, understand your users' sentiments, and interact with your users in a more casual way than you otherwise would with traditional software products.

Finally, all the data that you're able to gather doesn't just make for a better product experience and better user engagement, but it also makes it easier to predict future user needs and interactions. This makes the design process itself more fluid. It's not just that the ML algorithms are always learning, but so are you as a PM. Understanding historic trends in behavior and preferences offers PMs and designers a way to even further entrench their products into the psyche of their end users in a way that unequivocally demonstrates that they get what their end users care about and want from their product.

Machine learning

Sure, this one seems pretty obvious. We've gone over the requirements for the ongoing product management of AI products over various chapters in this book but it's worth echoing in this section as well. Unlike traditional software products, AI/ML products are not a set-it-and-forget-it kind of product. You're not establishing a list of rules that you'll update and check once a quarter or using a static data sample to power your product (most likely). You're creating an AI system that will support the ongoing maintenance and evolution of your product overall.

This includes data and ML models you're using to power your foundational product, data and ML models that will handle incoming activities, as well as data and ML models that will handle outgoing activities related to your product. That's a lot of data and modeling to keep a handle on, which means your product will evolve over time. Data drift, model drift, and evolving customer needs and preferences are all factors that will necessitate changes to the AI/ML algorithms that will support your AI products, particularly as you scale. Remember from earlier chapters that your data strategy and management will be an integral part of your AI PM work because AI systems require large amounts of data to get running optimally. With your product live and in production, your data needs will increase, not decrease.

Your product design team will need to account for this flux, whatever it may be. It will be hard to know upfront, of course, what that flux will look like. But there will be flux nonetheless. Training, adjusting, hyper-parameter tuning, and retraining your AI/ML algorithms will be part of the continuous maintenance and refinement that your product will undergo to make sure your models are up to date and accurate, working optimally and aligning with evolving customer needs. Building that time and resource allocation into your product design flow will be a key part of your AI product strategy.

Often, AI products will also be accompanied by features that offer a dashboard of trends, insights, and recommendations. Translating all that data and model activity into a framework that offers those outputs in the form of informative action steps for your customers is a moving target, but there will have to be a way to account for all the adjustments in data/model activity and to then package those adjustments in the form of recommended next steps for your end users. In traditional software products, your requirements might be more limited and externally set. You might do some market research, talking to potential users to arrive at a list of set specifications for your product, and they might routinely change based on ongoing efforts. But with AI products, this work is constant and incessant. If you don't have a plan in place to capture this complexity, your product and customer success teams will be overwhelmed.

Explainability

The consumer is changing, particularly when it comes to AI and the role generative AI has played in making AI capabilities accessible to the general public. The marked shift that came with the release of ChatGPT has implications for not just how consumers use and purchase products but also how they advocate for and evangelize those products. Whether you're on the B2B or B2B side of PM, you're going to have to do some explaining to somebody, and how you're going to do that will be wrapped into your brand strategy whether you like it or not. This is because AI ethics are becoming woven into a company's brand identity. Are you going to design the kind of product that's created by a company that's committed to data privacy and AI sustainability… or one that's not?

Consumers might not yet be well versed enough in AI to understand the variety of models that power the AI products they've come to know and love as household names. But that day is coming, and books like this one help with democratizing that kind of AI knowledge. Traditional software products don't need to worry about that kind of transparency as much as AI products do because they're scrutinized more heavily. Perhaps discussions around data privacy have accelerated with the rise of AI integration in products because AI systems require so much data and a lot of damage can be done with that data. Building trust with users is also important as data breaches and cybersecurity attacks become more common. Many AI systems will be incredibly vulnerable to these kinds of attacks.

Even if consumers don't have a deeper knowledge of AI, they do know that algorithms account for data points in a variety of weights, so sharing more about which models give more weight to which factors will be valuable information for them to have. B2C consumers may be concerned about these things for their own personal use, safety, and responsibility. B2B consumers will want to understand how your product works before their companies invest in you. They also may be subject to fines or litigation if they don't do their due diligence to properly vet your product.

You don't have to divulge a lot of information to offer your end users a glimpse into the inner workings of your product. Particularly with neural network-powered generative AI models, where determinations might not be so obvious for the builders of the products, there has to be some way to offer insight into how AI products are making certain decisions. Automated decision-making and determination can be scary and inaccurate in the best of circumstances, when you have tons of clean, representative data and you're building your product with integrity.

Many consumers have seen what happens when that isn't the case. Certain AI products may have frozen them out of potential jobs or mortgage rates if they haven't been de-biased. As a PM, the most encouraging part of the cultural shift as it pertains to AI products comes from bias. Consumers don't want to trust AI products that are biased and they reward brands that make a commitment to transparently proving that they're building ethically in any way they can. Soon, we will start to see more regulations that demand this kind of auditing and it won't just be largely consumer-driven. Building thought leadership, communication, and documentation that demonstrates your product is fair, unbiased, and accountable should be a big part of your AI product design. You'll be rewarded by your customers, peers, end users, and critics alike for it.

The design process needs to account for all the challenges and opportunities AI brings to a product. You're not just designing a product that will be used by a user. You're designing a universe that will be experienced by a variety of cross-functional teams, stakeholders, customers, and observers. Your product will need to communicate to the world that it understands why it exists through marketing, customer/user engagement, and, frankly, success. It will also need to communicate that you're using AI/ML responsibly and ethically.

> One of the central ideas we keep coming back to when we talk about AI products, as we have done across many of the chapters of this book, is that AI adds more complexity to an already complex ecosystem of jobs-to-be-done. Product management is already a multidisciplinary, complex role. You already have to balance a lot of hats. You're zooming out and assessing how your product is doing from a high level and you're zooming in and thinking about which story points need to make it in the next sprint. You're working with a design, development, and go-to-market team. That's a lot of variance for one role. It's a constant juggling act that needs to be done by someone who intuitively and deeply understands the product your company is building. That's still the case with AI products, but as we've just seen, you're balancing even more factors than you otherwise would with traditional software products.

Choosing your priorities wisely

We started this chapter with an overview of general product design principles and some of the top factors that impact the AI-native product design process. Now, let's turn to the essential task of setting priorities, which lies at the heart of the AI PM role. The work of product management always eventually turns to discussions of priority. *What do we work on first and why*? That prioritization will not only be informed by your product strategy but also by your leadership's values and the company mission as well. The product function doesn't live in a silo. It's an extension of the leadership at the company and the overarching values of the organization itself. Prioritization is impacted by organizational values, the broader mission, and the interplay of competing needs.

As you're establishing priorities for your AI-native product, it's important that you're maintaining product clarity and a shrewd balance of complexity, and aligning this with your product and company's greater mission and brand. We will cover these topics in the upcoming sections, but first, let's explore some foundational considerations for prioritization:

- **Purpose-driven priorities**: Make sure you're building teams and working in organizations that you believe in. At the individual level, we all have our own priorities for joining certain missions. Within the confines of this book, we talk a lot about product roles and the way things ought to be, understanding full well that we don't always have the luxury of doing things in order and with full resourcing. Limitations like resource constraints and time crunches exist. We might not always have the perfect leadership, team, or budget, but we can always decide if signing on to work on a product or with an organization is aligned with our values before committing.
- **Trust in leadership and teams**: You need to trust your leadership in order to manage their product. If you're a CEO or technologist, you need to trust the people working on your product in order to be effective. So, before we get into priorities and how to balance them, make sure you work with people who you feel have the right balance of values and integrity so that you can do your best work. Irrespective of the role or whether you work in the product team, it's a good first priority to have.
- **Communication**: Now that you're working with a functional team and your values align, you're going to be influencing quite a lot with the decisions you make surrounding product design priorities. There are multiple considerations, such as the AI architecture that will support your AI product, the aesthetic look and feel, your customer journey, and the features that will be included in your product. Prioritizing these aspects in a way that's purposeful will be the guiding light that will keep you grounded should you get off course. Communicating your priority list during your product strategy meetings will be integral to making sure you have buy-in and alignment from your leadership team.

Now, let's look at a few considerations about clarity, complexity, and branding that will inform how you prioritize your product design elements further.

Ensuring clarity

It can be easy to get swayed by the feature farm. This is a phenomenon that happens quite a lot with product managers, where they're getting so much feedback and so many ideas from the well-intentioned advisors and customers they have around them that they suddenly have a long list of backlog items to eventually include in their product design. The role of AI PM is so demanding and all-encompassing that psychologically, it feels good to have such a list. It gives one the impression of a kind of intuitive roadmap of where to go. Emotionally, it makes the PM feel like they are moving in the right direction and that the answers are all around them. But, often, all this does is add more confusion.

Ensuring clarity through the product design process is essential to creating a user-centered, intuitive, and effective AI product. Let's look at the importance of clarity in each stage of the AI product design process:

- **Identifying the end user**: Distinguish between user needs, pain points, and desires so that you're bringing in the right foundational elements for your AI product to address. Map out how your user personas interact with these needs, pain points, and desires and remember that there isn't just one "end user." You will be crafting a solution with the needs of multiple types of users in mind. Build distinct characteristics, behaviors, and needs for each persona so that you're capturing the full scale of the end user's needs. Build user stories with common user scenarios early so that you're clear on all the user personas you're designing for.
- **Defining the problem**: It's important to be clear about what you want for your product and what it ultimately exists for. You're working on an AI product and there are no doubts that you can include a lot of bells and whistles, but should you if they come at the expense of your product's intuitive UI and ease of use? Product management is just as much an exercise of finding what to include as it is deciding what to exclude. As AI PMs, we need to get comfortable with disappointing people and saying no, or at the very least saying "not now." Get clear on articulating the problem you're solving in a single, clear sentence and limit external influences that distract from that wherever you can.
- **Experimentation**: In this phase, the only thing that needs to be clarified is that this is the phase where all honest perspectives are welcome. Learn about ideation and brainstorming techniques that help encourage creativity and a diverse range of ideas. Try to ensure that all voices in the team are heard during these sessions. These ideas will translate to product features, functionality, wireframes, and prototypes to work with and show to prospective end users or beta testers. As you get more and more great ideas coming in, you'll be able to validate them to assess which are most feasible and most aligned with user needs, but in this phase, let your team experiment without fear or criticism.

- **Validation**: Now that you've been able to dip into the unknown, creative space to experiment, you can take stock of what has come out of that. Get clear on which ideas address the core problems discovered in the first phase. Validation is about reducing unnecessary complexity and functionality that has arisen in the experimentation phase. From the many potential ideas, you choose the best based on feedback. Put your prototypes in front of actual users and get their perspective on usability, intuitiveness, and functionality. If you came out of the experimentation phase with three potential prototypes, you'll use the validation phase to settle on one. Perhaps elements of one prototype resonated well with certain users over others. Perhaps some combination of all three is what makes it to the iteration stage. The prototype you end up with will reflect the core problems of your end users and will continue on to the final stages.
- **Iteration**: Validation is an ongoing process. It's formalized in this list as a step of product design but it functionally never ends. Even post-launch, when you're in the development cycle, you're always iterating on your product with each new release. The iteration stage can be difficult because it can introduce scope creep. As you're validating with more users in the prior phase, you might be tempted to add more features and functionality to the prototype to address all user concerns and issues. But remember, you have an entire product development cycle to worry about that. Here, you're getting clear on when your product design prototype is "good" enough to warrant branding and aesthetic final touches so that you can continue on to the work of developing and launching.
- **Aesthetics**: By this point, you've been working with a prototype that addresses the main issues and concerns of your end users and you've likely done a few rounds of iteration to make sure the inner workings of your product prototype are serving needs as intended. You've defined what the AI product is built for and that it is effective enough to visually design. At this point, you're getting clear on design elements that you need to include to make sure it's intuitive and easy to use and aligns with the overall company branding, identity, and values. Visual elements like fonts and colors, images, shapes, lines, borders, layouts, whitespace, patterns, interactive elements, hover effects, charts, graphs, infographics, page transitions, and logos all have a role to play in telling a story to your end user. Make sure that this story is as clear and concise as you can make it.
- **Documentation**: Now that you've gotten clarity on every stage of the design process as an AI PM, you have a pretty good sense of the narrative you're telling your end users and internal stakeholders. Documentation is about making sure that this narrative is as clear and honest as it possibly can be to the outside world. Documentation should be crafted from the design stage because design decisions, rationale, and specifications impact product decisions downstream once you are in the production cycle. Developers need to have as much information as possible when they start to turn the prototype design into a fully fledged product.

A lot can go wrong if you're not able to maintain clarity and objectivity through these various phases of your product design lifecycle. This explains why user research and UX are such important disciplines. They bring objectivity to a process that can get lost in its own subjectivity. A reality check is helpful when you're investing significant time, energy, money, and resources into a project. False assumptions can be costly and may convince you to go to market with a product that no one asked for and no one likes. Or you could launch a product with so many unnecessary features that your end users get overwhelmed and confused about what your product actually does.

Take the time to invest in research and inquiry. Understand where your users are emotionally with regard to their pain points and their preferences. Why do they like the things they want? Why do they dislike the things they don't want? Then, from there, how can you design an experience that meets those needs and preferences? It's better to start with a simple, elegant design than it is to start with more features than anyone knows what to do with. Less is more.

Getting clear on what your product actually needs to be will help maintain the experiences and contexts that your end users are working with. Spending time on research will also make your leadership feel better about the significant resources they're about to deploy on investing in an AI product. Of course, you can validate things as you go. It's best practice to have product designers, UX designers, and UX researchers as part of your product team full time. A lot goes into the work of capturing the evolving needs of an AI product. But make sure you spend time on gaining clarity upfront and not just after you've launched an MVP. Starting from a flawed product design will have you playing catch up and will ultimately prove to be more costly down the line.

Adding complexity

Eventually, you will build and add more features to your product MVP and this will set off another exercise in prioritization: *which features will come first*? If you're managing a B2B AI product, you might be tempted to go with the features that will satisfy your marquee customer best. If you're managing a B2C product, you might go for the most requested feature from your user feedback interviews. Both could be flawed choices because they only represent the thoughts and opinions of very select, outspoken, and privileged perspectives. The choices you have at given periods of time will have downstream impacts on customers you haven't acquired yet, that is, customers that you will acquire in the future.

The decision of where to add complexity and how shouldn't be taken lightly. As you start your validation stage, you're going to be confronted with more complexity with each new user or type of user you ask for feedback. Then, each time you dip back in to iterate and validate the next iteration, you'll see your list of considerations, features, and functionalities start to grow. You're going between validation and iteration several times before arriving at a final prototype. But you can't consider everyone's desires and preferences. The reality is you only have so many resources to make one "great" prototype, so complexity has to be limited.

Regardless of the type of product and market you're working in, once you have that one "great" prototype, the users of it will have to learn how to use it. There will be a learning curve of some sort and every degree of complexity you add to your product will have to have its own onboarding lesson accompanying it. Products that are cognitively heavy to process and understand will be faced with the most user resistance. But you can explore strategies to counteract this resistance.

Let's look at some important aspects that an AI PM must keep in mind:

- **Cognitive weight**: Remember that your end result is the prototype your development team will work on building into a fully fledged product for end users to use. Even if your prototype could theoretically capture all user needs and preferences, the end result might be clunkier than the average user's needs. This means they might have to sift through too many features or functionality that don't help or apply to them, which could lead to them getting frustrated and stopping using it altogether. The harder your product is to understand and use, the more frustration it will create with end users.
- **User adoption**: Many companies rely on having community managers or developer advocates in addition to customer success folks who help with keeping their community of users properly supported in cases where they have complex products. Often, this can be a source of bonding for the community of worldwide users to create a subculture around their product's brand identity. The issue of complexity doesn't necessarily have to be a challenge. More complexity can mean more capabilities and more nuance; complex products don't necessarily need to signal difficult products. Even if your product is challenging, and the road to onboarding and training on your product is a long one, you can still maintain a sprawling, bustling, loyal customer base. Good things may not come easily and the difficulty of mastering a product could be the very thing that fosters long-term adoption and success.
- **Cost**: You'll also have to balance the organization's appetite for additional complexity, particularly when it has to do with more AI features. The decision to store, use, and train on more data or introduce additional models to your product's design will also have to be handled collaboratively as more complexity in AI infrastructure can be costly, particularly if there are issues that might impact the scalability of your product design. Decisions like this will need to be strategic ones, and that additional complexity needs to be a bigger value addition than just adding more nice-to-haves to your product for the sake of a competitive edge.
- **Risk impact**: Additional complexity also means there are more opportunities for your AI product to get things wrong. We touched on ethical considerations for AI earlier in this chapter, and that's important to note here again as well, particularly when that added complexity is another AI feature.

The priority of added complexity is one that will have to be balanced with the customer, cost, and risk impact. You may find yourself having to answer questions such as:

- Will the added complexity hinder bias efforts in the AI system that supports your product design?
- Will it introduce new risks to determinations that might limit or restrict protected groups of people/users?
- Is there a way to test for that before you roll out that feature and see its impact on your real customer base?
- Is the impact of that added complexity fair across your entire customer base?
- Does the added complexity introduce data privacy concerns?

All these considerations have a role to play, and prioritizing ethics and risk when you're introducing more complexity into your product will help you build ethically and responsibly.

Branding

It might seem strange to see branding included in the list of priority factors, but branding is all around us. We all have a personal brand, whether we're aware of it or not. Our departments and teams have a brand identity, as do the companies we work for. When you're prioritizing the most important factors for your product design, you will have to think about how branding plays a role and if that branding aligns with the choices you're making at the product design level in the aesthetics phase. On a foundational level, the product you're designing must adhere to the brand identity the company has spent time creating.

This is because branding affects everything, from influencing your users and customers, building trust with them, and differentiating your value in the market, to emotionally engaging with your audience, your perception in the market as a company and as a product, and the role you play in ethics and social responsibility. When we're building and shipping products, many intuitive decisions are supported by the analytics and metrics we use to guide us. We implement product KPIs and metrics to let us know when we're on the right track.

In the absence of formal metrics, we can often revert to **brand identity** to be our guide. There will be emotional elements to branding and, as we build, these elements will shine through our products whether we know it or not. Brand identity is the vehicle with which we will communicate the company and product's values to end users and customers.

Even if you do everything right, how your product fits into your brand identity still might be a challenge. You've built an impeccable design, you've validated all your hunches, you've brought the best minds at the company on board, you feel your product is cohesive enough that its purpose and value are clear...and your product still might not be aligned with your brand identity. Particularly today, when AI products are ever present, aligning products to your brand and values is an important way to emotionally connect with the people who will buy and use them.

Branding isn't that nebulous a concept when you think about it. It's a voice, a tone, certain colors, and certain values and characters. It helps make your company and offering special in a sea of generic gray matter. As AI continues to be adopted, more and more companies will give in to the temptation to use generative AI for their marketing and communication efforts themselves. As convenient and cheap as outsourcing creativity to AI is, eventually it will make that sea of gray matter bigger. So, remember that as we continue on the AI journey together, showing more of your personality, identity, and values will be crucial because you're not selling to AI; you're selling to human beings.

We've discussed the concept of building trust as it relates to AI ethics and data privacy, but we need to come back to that idea here as well. Establishing a strong, reputable, and distinguishable voice through your branding efforts will instill trust and credibility in your AI product. Many AI products already exist in all verticals and industries, and the space will continue to get saturated as the decade continues. Differentiating your company and product intentionally through the way you communicate with your community will be key.

With more and more AI products crowding the space, you'll find yourself coming back to the same question when making product design decisions. "*Does this product design decision align with our brand identity?*" If upholding ethical standards, avoiding bias and discrimination, respecting your users' privacy, and promoting transparency and accountability are part of your brand identity, you'll want to signal that in various ways through your product design and communication efforts.

Product design decisions impact everything from how you'll set up your infrastructure and build your teams to how you'll craft a user experience, which ethical considerations you'll prioritize, and how you want to emotionally engage with your end users and customers. These will all be discussions to be had with all the key stakeholders you identify when you start the process of product design, and they're big discussions because they impact how your product will be perceived and engaged with. A product that is well designed, works well, and is valuable will have a good reputation in the market. That will be especially true if it's made by a company that's established and respected. It will lead to organic growth from users who are your champions, which will drive demand and keep your adoption high. Your UI, languages, overall UX design, colors, visual elements, design elements, and workflows are all part of crafting a visual brand identity in your product design. These are the tools you will use to create an experience for your end users and customers that instills trust and loyalty for your product over the long term.

I won't tell you which priority is correct; no book can. That's a collaborative effort that only you and your team can agree on together. But we already have a blueprint to work off of. Based on our product design principles from the first section, you've already got a place to start: your end user and market. Beyond that, we've also discussed the factors that add complexity to the product design process when AI is involved and some of the considerations you can entertain when you are prioritizing your product design efforts. These are meant to be tools you can deploy when you're starting out with bringing an AI-native product to market and initially setting out on the design journey. But that journey will continue, and if it's to be a long one, you'll need to communicate to your community about where you started and why.

Think about that today. Why has your company already invested in getting your own team resourced? Surely there must be a compelling reason your position at the company exists or why your company decided to go all in with launching this new AI product. Getting down to the bottom of that narrative and finding the passion that has led to those decisions is the kernel that will make up your story. Communicating that story will be a big reason people buy your product.

What's the story you're telling?

Everyone loves a story of tech founders who believed so much in their products that they built their companies in their garages. Often these stories are false or flawed, projecting an image of a single founder who rose through adversity through their passion, dedication, and bootstrapping alone. Eventually, they get lucky and the right person sees the value in what they're doing and, inspired, they go all in and bankroll their noble ambitions. This is a kind of myth, of course, and we humans have been interested in mythology for millennia.

Storytelling is like candy for our brains; it helps us cut through all the distractions in our lives to intrinsically enjoy something. That enjoyment allows us to do a number of things. Firstly, it can help us better understand complicated things like AI products with a steep learning curve. Also, engaging users and communicating why your product matters in the first place is how you can start to build a relationship with users, and that's taken to the next level when you're able to use that engagement to make your product more digestible and add more context to how it works and why it was built the way it was.

Because you're creating a narrative around your product, this also makes your offering more personable and relatable. According to a Nielsen study, people buy products largely for emotional reasons. Whether you're selling in B2B or B2C environments, storytelling offers a cheat code to effectively communicate the worth and emotional benefits of purchasing and using our products. As we discussed in earlier sections of this chapter, you will be finding ways to emotionally resonate with your end users through the elements of your product design. But the storytelling that will accompany your product will also be doing that. Having and building on a compelling story will make it so that others will truly empathize with what your product is trying to do on multiple levels.

There are also highly practical reasons for investing in storytelling. Through the act of listening to a story, we're able to remember quite a lot more detail than we otherwise would if our emotions weren't engaged. Some scientists have discovered that our brain releases chemicals like cortisol, dopamine and oxytocin when we're being told a story. These chemicals help with formulating memories, regulating our emotional responses to keep us engaged, and enhancing our empathy to help develop and maintain relationships.

Particularly in the context of AI products, where multiple use cases may be involved, building on storytelling is a way for us to ensure that customers not only understand our products but can see them playing out in multiple contexts. Not only that, but they'll find those varying contexts memorable as well. We've discussed the ethics of maintaining AI products throughout this chapter, and storytelling is also an excellent way to create narratives around why this matters to your company, brand, and product. Doing so will make sure that your brand perception is tied to your commitment to building AI ethically and responsibly.

Storytelling can be a number of things. It can be your CEO or leadership team evangelizing the work your company is doing at conferences and events. It can be in the form of blog posts, white papers, special reports, or content on your company website/product page. It can be in the way you comment and engage with your online community. It can be in your own social media strategy, language, and posts. Or it can be baked into the product design and user experience itself.

Now that we've talked about the benefits and dynamics involved in storytelling, let's see them applied across the product design itself. Incorporating storytelling elements will help with a number of things: communicating the value and relevance of your product, engaging your user and customer base, and offering an emotional connection to those who are engaging with your product. Building a narrative that brings your customers and users closer and keeps them engaged will ensure that you nurture a loyal following for your AI product. Let's explore what storytelling elements might look like when they're embedded into your product design efforts. One of my favorite products to use is Duolingo, and I will use the app as a positive example of how a company can use storytelling in the delivery of their AI product.

Set the stage

Duolingo is consistently creating narratives around why the product exists and how it helps people. I know why I'm using Duolingo, but I don't necessarily know why or how others are. Setting the stage is important because, experientially, it constantly reinforces why something is relevant to me. This kind of framing consistently helps me understand the importance of the *value* of the app. As I continue on the app, I am reminded often about how Duolingo is effective for others, which percentile I'm in when I complete a task or milestone, and when I've hit a pain point or a celebratory moment. Context is key, and while these narratives are not too overwhelming or constant, they consistently remind me of my orientation within the product.

Characters

We see this with AI products all the time. Whether it's a chatbot feature or a voice, some sort of character is involved, which anthropomorphizes the AI. In Duolingo, there are a number of characters of a variety of ages that come out to interact with me, and each brings a different energy or perspective to the experience of the product. The diversity of the characters also plays a role in the gamification of the app, which is already intrinsically designed to make me want to win…and learn French as a consolation prize. The characters also allow me to feel like I am indeed playing with a team, and that there are a number of other entities that are there to see me succeed. Yes, I can suspend my disbelief enough to feed into the fantasy if it means getting to the next level.

Progression

Every good story has a compelling intro, a slow build to a climax, and a resolution. You can imagine the unfolding of your product in much the same way as you build your design. Perhaps you don't want to unveil all the features to your users at the same time so as not to inundate them with capabilities. Perhaps achievements of some kind can unlock other features. If there can be some progression that your customers and users experience, it will give off the effect that they're on a journey with you and that their actions and performance during this journey matter, that they have real consequences. Preventing overwhelming cognitive load is tricky with all tech products, including AI products.

In our Duolingo example, the app does this expertly by showing me exactly what steps I have coming up in my journey. I can see all my learning courses outlined as a step-by-step roadmap. I can decide how fast I want to go and see what's coming up as I advance. Bringing the users on a journey like this also contributes to periods where they are delighted and surprised by what they discover or unlock. That boosts engagement as well. The progressive quality of storytelling in your product design process will also capture a user's milestones as they continue their journey. Building narrative components for when your users hit a goal, meet a requirement, celebrate the completion of something, hit periods of difficulty, or are given feedback to help them improve are all ways you can demonstrate the progression they're experiencing as they continue to use your product. It also communicates that you understand them as a user and as a person.

Knowledge

Building a strong knowledge base that communicates the intended use and special perks of your product experience is also key to boosting engagement. Particularly early on, when your customers and users are onboarding or when they become part of a bustling community, you'll want there to be enough material about your product design, UI, and UX to keep those who are knowledge-hungry fed. Keeping a robust knowledge base is a way of keeping the storytelling alive with your user base.

In the case of Duolingo, there are a number of articles, resources, and sections that cover various areas of learning a language, and they make a point of alerting me when there is something new on the app. We touched on documentation earlier on in the chapter and this is an extension of that. All customer-facing resources and documentation are opportunities for you to announce new features, highlight core offerings and concepts, and reinforce important interactions that you want your users to be having. Having tutorials or onboarding materials that can be ingested as stories will enable your users to learn and remember and reinforce why your brand and product matter to them.

Call to action

Often, AI products include some sort of call to action at the end. Perhaps it's a recommendation, a next step, or an explanation of some deterministic quality it's made. You can use narration when delivering a call to action to your customers to give them more context and better explain why your AI has come to that decision. This does two things:

- It makes AI-based determinants more human and understandable.
- It offers your customers and users more of a reason to believe in your product.

Because building trust is so important with AI products, particularly now when fears of AI going wrong are still alive and well with so many, building narration into all aspects of communication with your customers and users is a way to reinforce that trust you've spent so much time building! Duolingo offers a number of ways to spring you into action. They might have you on a friend quest, so you're galvanized by beating a friendly adversary. If you're making certain mistakes often, they give you a brief window of time to understand a concept before giving you the next lesson. If you are progressing well, they also alert you to potential new awards and badges you can receive if you keep up the good work. All these are ways where AI features and mechanisms powering an app can be used to influence behavior.

Case study

Using our case study example of Waterbear, we see a number of these storytelling elements alive in much the same way as Duolingo. The Waterbear founding team wanted to offer an experience that was gamified for users, particularly because they felt this would help balance some of the heavier aspects of using a mental health app. Though gamification works and can promote engagement, as we've seen with Duolingo, it also has the capacity to detract from an app experience. The founding team found early on that their first iteration of the product actually wasn't resonating well with end users because they felt the gamification was too lighthearted for what the app was intended to do. The feedback was overwhelmingly that although they appreciated that the app "encouraged" you as if you were playing a game, it felt tone deaf. Users were, after all, not playing a game but instead sharing deep, emotional, and sometimes overwhelming topics with a conversational AI that was trying to get to know them. This was one of many product design lessons learned for Waterbear.

Let's explore some of the lessons learned based on the product design elements from earlier in the chapter:

1. Start with understanding the end user—really!

   Initially, the Waterbear founding team had lofty dreams of completely outsourcing mental healthcare to AI. AI was the silver bullet and they overidentified with the hype! They thought their product could solve the issue of there not being enough therapists in the world and the high demand for support for mental health needs. But it turns out that replacing a human therapist isn't really all that easy. There are a number of therapeutic disciplines a professional can build a career in and one model couldn't possibly account for them all without contradicting itself. Eventually, they started talking to real potential end users of Akeira and found something quite surprising, which allowed them to define the problem more easily.

2. Define the problem before it defines you.

   What they discovered was that none of the potential end users actually wanted to replace their therapists, or use Akeira in place of a therapist. All they wanted out of Akeira was a friendly ear and a place to centralize all their mental activities. They didn't want to be diagnosed or psychologized. They just wanted a safe place to share. One of the resounding pain points was that it was hard to have a consistent place to just voice mental health concerns and try to make sense of them.

3. Experimentation: Start with a few prototypes and then settle on one.

   Defining the problem through a more user-centric lens helped inform the Akeira product team on the kinds of solutions they could experiment with and validate before arriving at a final prototype. But they settled on an Akeira prototype too quickly, without bringing in enough diverse perspectives and considering at least two or three prototypes. The first prototype had a journaling function, a dashboard function, and a chat function. Given this level of complexity for the first prototype, it was hard for the product team to get consistent feedback from beta testers.

Here is an example of a wireframe Akeira's product team produced to build the final prototype:

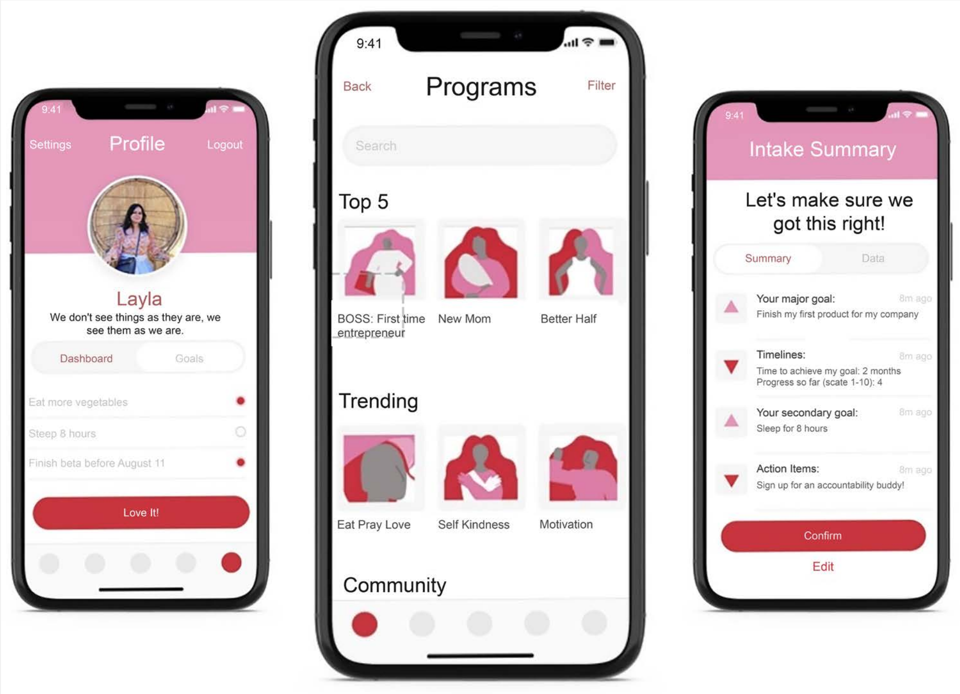

Figure 9.2: Reference wireframes for Akeira

4. Validation: Testing assumptions.

   Some users found the chat function really helpful but the diary function awkward. Some liked the dashboard but others found it confusing because they felt it wasn't clear or deterministic enough. One of the most resounding pieces of feedback from Akeira's beta testing group was that they wanted Akeira to serve as a companion tool they could share with a therapist or a loved one if they wanted to open up. This was still pretty early on and the company still did OK along its product design journey, but it's a great reminder of what can happen when certain assumptions are made too early on before proper vetting from real end users.

5. Iteration is key.

   When it set out to launch its flagship product Akeira, the leadership and product team at Waterbear knew that they wanted to go to market with a mental health-focused app, but early on they had initially built an MVP in mind with the ChatGPT API. What they found after a few different iterations of the product using other foundational models by other companies (LLMs) was that performance was improved by the BERT model compared to ChatGPT's model.

They found more hallucinations with ChatGPT than they did with BERT and decided to go for the more conservative model given the subject matter was mental health focused. They didn't want the model to say anything sensationalizing or triggering to someone who was being vulnerable with their technology and trusting their product design choices.

6. Aesthetics don't have to be decided internally.

   The Akeira product team also wanted to include some of their product decisions on their blog so that the loyal community they were building was let in on some of these decisions. The choice to do so was twofold: they could show off their company's transparency as it relates to AI and it would also allow them to get an early signal from eager customers who wanted to be a foundational part of the building process. It turned out that their online community appreciated the transparency so much that they started recruiting more potential customers by sharing the insights on social media. Akeira was able to double its **monthly active users** (**MAUs**) within three months following that blog post.

   Partly, it came from customers feeling let in on a secret. They felt trusted and appreciated enough by receiving a special blog post about it. But they also felt like it was an opportunity to learn. By posting that blog that day, they felt their intelligence was being complimented, not insulted. These were not end users who were particularly interested in AI, but the candor actually got them interested in not just AI but Waterbear itself. By identifying so much with the brand and feeling like they were let in on a secret garden of possibilities, they found themselves more loyal to the company. Other companies might not have chosen to be so transparent, but in the case of Waterbear, it had a galvanizing effect especially because it is so rare.

7. Documentation: Including your customers in documentation and product design choices can pay off.

   Part of the hype that came from the blog initiative mentioned earlier bled into potential new product areas for Waterbear. Eventually, there was enough discussion around the company and its product choices that potential customers outside of Akeira's demographic were asking for their own products. Akeira was built for and targeted women over the age of 18, but some of the online discourse came from men as well as girls under the age of 18. Waterbear hasn't launched any new products yet, but it has begun the research phase of its product design lifecycle to begin developing new products for girls and men. Preliminary discussions have begun for boys under the age of 18 as well, but no product design activities have begun in that area yet.

In the end, the Akeira product team found a way to beautifully balance the unknown with the constants they reassuringly came back to. They knew from the beginning that they wanted to build a product experience around real value and wellness for their users. They also had the humility to keep coming back for feedback and iteration until they felt they had accomplished that.

Summary

In this chapter, we covered the various stages of product design to cover the foundational design elements that will inform how an AI product will be conceptualized and brought to market through the development of a prototype. We also covered the areas that make AI-native product design especially challenging, focusing on the specifics of bringing forth solutions that are user-centric, as well as the ML and explainability challenges. We discussed what AI PMs can do to ensure they stay clear and consistent with their objectives as they're going through the product design process and finished with a discussion on storytelling and the role it plays in influencing product design. These are all relevant factors that will play a vital role in ensuring your nascent AI product's success and adoption because they help communicate who your product is for, as well as the story of how it will help them along their journey.

Product design is about building a blueprint that will grow into a fully developed, managed, and mature product over time. This was a dense chapter that set the stage for the next phase post-prototype launch: establishing performance metrics for your prototype. These are metrics you will be incorporating once you're ready to build your fully fledged product and enter the development cycle.

References

- *The neuroscience of storytelling*: `https://neuroleadership.com/your-brain-at-work/the-neuroscience-of-storytelling/`
- *When emotions give a lift to advertising*: `https://www.nielsen.com/insights/2017/when-emotions-give-a-lift-to-advertising/`
- *The science behind the art of storytelling*: `https://www.harvardbusiness.org/the-science-behind-the-art-of-storytelling/`

10

Benchmarking Performance, Growth Hacking, and Cost

In this chapter, we will understand the benchmarking needed to gauge product success in all its various forms. We will be looking at performance metrics from a product and growth perspective. These include **value metrics**, **north star metrics**, **key performance indicators** (**KPIs**), and **objectives and key results** (**OKRs**) that companies can use to get early signals on whether their product strategy is successful and will lead to growth hacking. We will also be discussing how companies can prepare themselves to defend their product when it's compared to other products.

When it comes to the success of your product, it will all come down to how effectively you can communicate its value to your potential customers and put it to the test. Part of maneuvering between your sales teams and customer conversations will involve making sure the tangible benefit your AI product can bring to your customers is a big part of the sales process. When customers come to you and give you the stage, do you have a dependable demo to show them? Are you able to show them your product's performance across a variety of datasets or does it only work well when you use your own data? Can you talk about the approach or algorithm you're using without having to resort to the safe space of a *proprietary blend* that you're not at liberty to discuss? Can you confidently get them in touch with partners or customers they can speak with to get a referral? Before you start to track the performance of your product for your sales or investment pitch, you will need to build trust within the ecosystem you're a part of. Communicating success when there isn't a foundation of strength will ring hollow.

As we continue with the chapter and discuss how we can gauge and communicate a product's success, as well as how we can determine an appropriate cost structure, we have to remember one foundational truth: at the end of the day, cost doesn't matter. Sure, it's important to remain in a palatable realm. Does the price align with your product's value proposition? Is it competitive compared to comparable solutions out there in the market? Some of your customers will be more cost-conscious than others, but at the end of the day, customers don't necessarily forgo powerful products because of the price tag. If it works and it works better than the other products out there, they will usually find a way to make the budget.

If you're working as an AI PM at a B2B company, you're well aware that buying software is tricky because purchasing decisions typically involve multiple stakeholders. Amid the **proofs of concept**, the budget discussions, countless meetings with potential vendors, time and dedication of your own technical resources to evaluate every tool in your environment, feature comparisons, managing internal stakeholders, and, ultimately, achieving alignment on the product the company believes is best for its use case, the entire process can be exhausting and frustrating. If you're an AI PM for a B2C company, you still have these hurdles, only with fewer business stakeholders to juggle. This offers PMs a huge opportunity to make their customers' lives much easier while making their products stand out from the rest.

Decision-making in B2B contexts can be lengthy and involve many layers of approval, making it more complicated than in B2C scenarios. AI PMs must focus on enhancing user experience, addressing customer pain points, and managing areas of friction as their products are evaluated before purchase. In order to do this effectively, they must define metrics that measure product success and effectiveness so that they're aware of how their products are being received, whether they're B2B or B2C products.

By establishing value metrics, AI PMs make sure the decisions they're making are objectively supported by data. This means they can better justify investments they're making in AI/ML capabilities and third-party tools, as well as addressing concerns around costs and demonstrate value to customers and users. If AI PMs can build a product that demonstrating immediate value – a product that can help expedite decisions and reduce friction for users – they can create a compelling use case that resonates deeply with customers and eases the path to alignment and purchase.

With that, let's dive into the chapter. We will be covering the following topics here:

- Value metrics – a guide to north star metrics, KPIs, and OKRs
- Hacking – product-led growth
- The tech stack – early signals
- Managing costs and pricing – AI is expensive

Value metrics – a guide to north star metrics, KPIs and OKRs

In this section, we will be discussing value metrics, KPIs, and OKRs through the lens of what's best kept inside the business and what's best communicated outward. All relationships are built on trust, and this includes the relationships that companies/PMs have with their customers and the market. Assuming you've found a way to build a reputable circle of trust around your product, it's now your job as PM to make sure this relationship continues to grow and evolve, ideally by incorporating feedback from your market.

A great way to do this is by communicating the value of your product and continuously building and referencing that growth and evolution. The way to do so is by using tools to track and confirm your product's progress. The key question is: *what will you choose to track and which metrics will you give particular attention to*? This is something that all companies struggle with getting right because it can be tempting to tie all pursuits to revenue and make that the top value metric, but that often doesn't give us the full picture.

Choosing metrics, particularly a north star metric, is a rigorous and meticulous activity, and it's best done in a collaborative environment. You want to have a diversity of voices present when you're making such foundational decisions. These are decisions that shouldn't be taken lightly because, in many ways, the establishment of these norms will have downstream ripple effects that touch virtually all your internal teams in one way or another.

All companies and organizations have a certain amount of entropy they're battling at all times. Entropy is the force that introduces disorder, uncertainty, or randomness to how business is getting done, and it's the central force we're building against when defining success and setting goals. We've covered the concept of drift in earlier chapters and that concept applies here. You can think of entropy as organizational drift. This is because as the days, weeks, and months go by, priorities can change. Sometimes there is a good reason for this. But if you're planning and strategizing well, you should resist ambient changes to make sure you're executing the product strategy you laid out.

Building OKRs at the product- or org-level functions as a tool against that entropy. It ensures that you're sticking to the values, beliefs, and purpose that you set out to address by building a company and product. It's advisable that companies have 3-4 product or organizational OKRs to make sure that alignment isn't disrupted. OKRs are strategic choices a company makes to prioritize certain outcomes over others. They should be reasonable, achievable targets that are aligned with the product strategy and reflect the goals and priorities of the company.

The KPIs we will focus on in this chapter mostly relate to your general product's KPIs and how they relate to your overall product strategy. Your product strategy is an extension of your product vision, establishing KPIs and metrics that reflect not just what your product's goals are for its users but also what its users find most valuable about it.

Finding a way to articulate this value and communicate it continuously in both internal and external contexts is perhaps the most important aspect of engaging your market because it influences how you see and hear your customers and, ultimately, how you speak to them. Let's dive in!

North star metrics

Choosing your product's north star metric is an important part of aligning your product strategy with the realization of a key goal. It's common for most products to be oriented toward one top metric that the company is using to track progress. Choosing the wrong one, or not choosing one at all, can have a discombobulating influence on how your teams are built, guided, and measured. It can also have the unfortunate unintended consequence of not giving you the full picture of what's happening on the customer side. This is particularly important for the nascent AI/ML-native product because you want to make sure you're getting an early signal on what's working and what isn't so that you can improve your product sooner.

Depending on the forces involved with your product, you'll be subject to certain influences. This could mean influence from investors and leadership, but it also means influence from the builders and users of your product. Of course, we rely on feedback to make product decisions. But putting the users and customer voice central to product strategy isn't a given. Many product teams will call themselves customer obsessed, but few follow this lip service in practice.

This is why choosing a north star metric is structurally tied to your product strategy and your values as a company. This can also have a disruptive effect on how you choose to measure the success of your product overall. This is because the north star metric serves as your primary KPI. It doesn't just unite your entire company around a common goal, but it also communicates how well the company as a whole is performing. Using it tells folks internally and externally whether things are going well overall (or not). If you're doing it right, you will have one metric to serve as an elegant indicator of revenue, performance, and engagement.

It might be tempting to take a *laissez-faire* approach. After all, a product has traditionally and often been measured on how much is built and shipped. However, establishing a north star metric is a strategic choice that benefits the product's function and influence within the company. This is especially true for AI companies because of the novelty of AI and the strategic importance it has in the market today. Establishing a north star metric is a way of demonstrating to the entire organization that the product team is foundationally aligned with the success of the business, and it reinforces the idea that the company is product-led and not sales-led.

The benefits of doing this are threefold:

- You're giving your organization as a whole alignment with your product team.
- You're also giving your product team a standard to stay accountable to and be measured against.
- These two systems working together create a feedback loop so that anyone at your organization can see, with a large degree of certainty, whether the ship is sailing the way it should be. This informs discussions that marketing, sales, logistics, finance, product, and customer success teams all have with each other, as well as within their own teams.

If you're a PM, particularly one pioneering a novel AI product, we highly suggest you push for the establishment of a north star metric. Chances are, if you're foundationally involved with building your product strategy and creating a product vision, you have some clarity already on what this metric will be because it aligns most closely with the ultimate value of your product. We suggest Growth Academy as a resource if you'd like more direction on how to establish a north star metric, but for now, here are some questions to get you going:

- What does my product intend to do?
- How do I know whether it carries out its intended function?
- Which activities result in the best experience with my product?
- What do my customers value most about my product?
- Is there a moment in my customer's journey that hooks them?
- What are some indicators of future revenue other than *sign-ups* or *purchases*?
- Is my goal to make customers more productive, capture their attention, or get them to buy something?
- Do I want my customers to take a specific action or a wide variety of actions?
- Is success measured by how often or how efficiently my customers rely on my product?

Every company will have a different methodology for arriving at a north star metric because it's so foundationally centered around the business model, which is why establishing a north star metric is part of your overall product strategy discussions. It will also be based on your ideal customer journey because you're optimizing for a certain moment that most benefits your customer's expectations and aligns that moment with what's most important to the company's success.

> Note
>
> You may want to have more than one north star metric if your business model and success make more sense with a multifaceted approach.

The following is a list of famous digital product companies and their north star metrics:

- *Airbnb* focuses on the number of nights booked.
- *Uber*'s ultimate goal is to boost ridership, so they focus on rides booked.
- *HubSpot* looks for entire teams to be actively engaged on their platform to ensure their software is depended on regularly, so they use weekly active teams.
- *Amazon* doesn't just look for total sales – they look for a pattern of behavior for all their users, so they use monthly purchases per user.
- *LinkedIn* might not be a social media tool you'll use daily, but if you're a working professional, you'll likely have a reason to come back on a monthly basis, so they focus on monthly active users.
- *BlueApron* is one of the few that have two north star metrics, so they focus on orders per customer as well as active subscribers for their subscription model.
- *Facebook* and *Instagram* want to create a daily dependency on your attention, so they focus on daily active users.
- *Netflix* ultimately tries to encourage people to view all their films and shows and they keep it simple by focusing on watch time.
- *Spotify* also has two north star metrics, conveying a focus on returning users with monthly active users, as well as the consumption of music on their platform with stream time.
- *Duolingo* has three north star metrics – daily active users, learn time, and success – reinforcing their desire for active learners to come back frequently and spend a significant amount of time learning and seeing the positive results of that dedication.
- *Quora* keeps it simple by focusing on the number of questions answered.

There are also a number of out-of-the-box north star metrics categories you can contemplate as you're sifting through your options, particularly if you're getting started with a novel product in your product suite or if you're just building a company. Again, be incredibly discerning about these metrics and put them to the test. Brainstorm with your go-to-market/sales/leadership team and workshop them. Try to arrive at a full destination assuming these metrics are in place. Ask yourself: what does our business look like a year, 5 years, or 10 years down the line if we make this our north star?

KPIs and other metrics

The previous section on north star metrics gave us some good examples of the variety of metrics we can track. This is an excellent way of grouping various KPIs you're going to track, besides whichever you settle on for your north star metric, because the goal of metrics is to give us quick flashes of insight into the strength and success of your product. However, these numbers don't stand alone. What good is it if your revenue is great but it's only coming from two customers? What good is it if your customers install your product but never use it actively?

Having a well-rounded assortment of KPIs is responsible. Doing your due diligence to explore the variety of ways your customers and users interact with your product will go a long way toward not only your product succeeding in delivering value to your customers but also educating your product organization on potential areas of improvement and expansion as far as product features are concerned. This is doubly true for an AI product, which may have fault lines that traditional software products don't have. Using KPIs to track several behaviors and performance benchmarks is a way for you to keep the feedback loop between your customers and your product alive.

> We refer to KPIs and metrics interchangeably in this chapter, but there is one crucial difference. A **metric** is more of a generalized term that applies to something that can be quantifiably measured. A KPI relates to something specific that can be tracked and measured against key business objectives. Not all your metrics will relate to your top-line business objectives, but all your KPIs will be metrics. The point of using these is to see where there needs to be new growth in the adoption of your product. Maybe your north star metric tracks well, but there are dips in your other KPIs. This is an indicator that there might be new product areas to focus on.

The use of KPIs and other metrics is meant to serve as an early warning sign to PMs and leaders about the strength, success, and usability of their product, but beyond that, they also are meant to paint a narrative about why something isn't going well. This is why the orchestration of your KPIs and metrics is such a thought-intensive process.

You want to use your metrics to help you to troubleshoot problems that arise, such as a bug or an issue with a page, but you also want to use them to help discover areas of the product experience that are suboptimal. You want to understand whether there is a problem with how your product is being built but also whether there is a temporary problem with how the product is being deployed. Using an assortment of metrics to gauge the overall product experience is a best practice for AI PMs. The following is a list of many common KPIs and metrics you can consider, along with descriptions:

- **Annual/monthly recurring revenue**: A measure of predictable revenues across a month or year
- **Revenue growth**: The percentage of revenue growth compared to the previous period
- **Paid users**: The number of all active subscribers/purchasers
- **The average revenue per user**: All revenue divided by the number of total users
- **Customer lifetime value (CLTV)**: The revenue generated by each customer acquired
- **Customer acquisition costs (CAC)**: The total costs associated with marketing, sales, and advertising for each customer acquired

- **Profit margin:** The percentage (ratio) of revenue compared to costs
- **Market share:** The percentage (ratio) of your company's total revenue compared to your industry's total revenue
- **Monthly/daily active users:** The number of active subscribers/purchasers over a period of time
- **Net promoter score (NPS):** A metric of customer loyalty
- **Customer satisfaction (CSAT):** A measure of customer sentiment
- **Customer retention rate:** The percentage of customers that stay with you over a period of time
- **Churn rate:** The percentage of customers that stop buying from you or using your service over a period of time
- **Messages sent/nights booked/rides booked:** You can insert any applicable action here, but this would be one particular action you're using as a marker for success, such as nights booked or rides booked
- **Sessions per user:** The number of times a user opens your product/app
- **Session duration:** How long each user spends using your product/app
- **User actions per session:** How active users are when they are in session
- **Average support resolution time:** How long your customers wait to get support when they have issues with your product/app

While many of the metrics listed here relate to both business models, B2B products have KPIs that are specific to the nature of long sales cycles, customer relationships, and recurring revenue. B2C products, in contrast, tend to have KPIs that are specific to customer behavior, transactions, and consumer engagement. So, before moving on, let's look at KPIs and metrics that are specific to B2B and B2C products:

- **B2B KPIs:**
    - **Sales cycle length:** The average time it takes to convert a lead into a paying customer
    - **Lead conversion rate:** The percentage of qualified leads that are converted into paying customers
    - **Upsell rates:** The percentage of existing customers purchasing premium features/products
    - **Cross-sell rates:** The percentage of existing customers purchasing additional features/products
    - **Pipeline velocity:** How quickly deals move through the sales pipeline
    - **Account expansion rate:** A measure of growth within existing customer accounts through additional licenses or expanded usage
- **B2C KPIs:**
- **Cart abandonment rate:** The percentage of users who add items to their shopping cart but don't complete the purchase
- **Repeat purchase rate:** The percentage of customers who return to make additional purchases
- **Customer acquisition channels:** Common breakdowns of where new customers are coming from (social media, mail marketing campaigns, organic search, etc).

- **Referral rate**: The percentage of new customers who come from existing customer referrals
- **App downloads**: Number of weekly or monthly app downloads and installations
- **Customer engagement metrics**: Measures of depth of user interactions with the product, app, or website for things like time on site and bounce rate

Establishing metrics that are specific and measurable is a given, but they also have to be highly relevant to be justified. Having a high number of unique visitors, sessions, or social media followers might look great for a while, but these markers don't actually help you in the long run because they aren't actionable. There's nowhere to go with that information if those unique visitors arrive accidentally or leave because something's wrong, if the sessions are accidents or clickbait, or if the social media followers are bots.

You also may have heard the term "vanity metrics" before, and this is a point of caution for leaders and PMs alike. **Vanity metrics** are metrics that are flattering to the company but actually have no real basis for communicating real success or performance. They also don't tell you what to improve on. Make sure any metrics you're using, north star or otherwise, all provide the company with important insights.

When you're looking to establish metrics and KPIs, keep in mind the following questions:

- How does this metric influence our business/product decisions?
- What does this metric actually tell us about our product/customer behavior?
- Does the data from this metric actually reflect reality?
- Is this metric tied to a part of a specific process we can improve once we start measuring it?

OKRs and product strategy

We introduced OKRs in *Chapter 6*. Let's jog your memory – the idea behind **OKRs** is to tie certain results back to the overarching goals of the company or team. Usually, there are organizational OKRs, but individual teams and even individual people can also have them. OKRs consist of one objective that is aligned with various key results that reflect the attainment and success of that objective. They arise from a clear, defined goal and are matched with a series of success criteria to constitute the completion of that goal. You look at what you want to achieve and then you set guidelines for knowing when you've achieved it.

All of this work of establishing a north star metric and relating it to other organizational metrics and KPIs serves your highest OKRs and, finally, your product strategy as a whole. All of this orchestration is part of the in-depth work that has to be done to align organizational growth with product success. Doing so not only reinforces the product team's influence on and strategic input into the company but also aligns the entire organization with a high level of focus and alignment.

Getting all teams committed and oriented toward certain organization-wide goals and tracking metrics that support those goals allow the company to be data-driven and also allow product teams to make data-driven decisions about how their products are used and how that impacts the overall bottom line. OKRs are meant to be ambitious, and they're meant to stretch your organization or team's capacity.

Creating a measurable, metrics-driven product strategy is optimal, particularly for an AI product, because you need as much visibility as possible into your customers' experience of your product so that you know what to improve and optimize. Everything we've discussed so far – your product vision, your north star metric, your supporting KPIs and metrics, as well as your OKRs – should all support your overarching product strategy.

> We love the graphic presented by ProductPlan because it highlights this relationship between product vision and strategy, as well as how that strategy is then applied within the various supporting teams that drive the product forward. According to the *Product Strategy* blog published on the ProductPlan website, product strategy is high level. It's a plan that describes what a business would like to accomplish with the product it's building. The product becomes a vehicle for the business's overarching goals. Because it's a plan, it should address the personas and customer segments the product will serve, how it will benefit those segments, as well as the company's goals for the product after development starts. The product strategy, which aligns with all the metrics we've covered in this chapter, influences how the product interacts with customers and internal teams within the business, how it's built relative to competing products, and how it will relate to macro trends, with the aim of establishing a product roadmap. You can think of the product strategy as the overall plan, and the tracking, measurement, and improvement of the metrics and OKRs as the concrete, tangible way that you will implement that plan.

In this section, we've covered a few of the key elements involved in building out a product strategy and how to measure the progress of that strategy. There are several options a company can take with this. The direction a company or product team takes is highly dependent on the organization's goals, business model, market strategy, competitive landscape, and overall market conditions.

Next, we will cover some of the elements of growth hacking so that an emerging AI/ML-native company or product can get early signals on whether its product is meeting its customers' expectations, which will help lead to product-market fit and commercial success.

Hacking – product-led growth

There is philosophical debate these days about whether companies should be marketing-led, sales-led, or product-led. Depending on where you are in the ecosystem of a company, you might have a bias for favoring your own designation. As PMs, we might have this bias as well, but if we take marketing- or sales-led growth to its natural conclusion, we might be faced with a predicament. Let's look at these different perspectives for a moment:

- **Marketing** crafts a message that might resonate beautifully with the audience, but might not actually align with what's being built. Marketing-led growth focuses on driving business growth through marketing efforts like advertising, content creation, and branding. Attracting, nurturing, and converting leads into customers are the main focus. While this strategy may lead to quick brand awareness and attract customers, acquiring those customers could be costly, and it can make the business reliant on marketing spend for growth.

- **Sales** may line up the ideal marquee customers for a company to take on, but they will be forced to reactively dictate what product should be built based on a handful of customer preferences. Sales-led growth focuses on building dedicated teams to engage with leads through direct sales outreach, sales enablement, and relationship-building through networking and partnerships. Though this is an effective strategy for high-priced complex sales cycles, it is resource- and cost-intensive.
- **Product-led growth** is about making the product itself the primary driver of acquisition, retention, and expansion by allowing users to experience the value of the product directly through freemium or trial models. The focus here is on empowering prospective customers and users through self-service, in-product experiences, product analytics, and customer feedback loops. Optimizing product experiences and investing in product development are significant associated costs, and this strategy might be too costly for products that need to be highly customized.

The role of a product team is multifaceted and multidisciplinary, and their involvement with marketing and sales serves as an extension of their own work. A product can't exist alone. If a product is built and no one hears about it or buys it, does it even exist? There has to be so much extensive work done on research, market analysis, and understanding the competitive landscape and the true needs of the market in order to ideate and build a product that stands a chance. Often, product management is at the nexus of these activities and is there from the very beginning.

That isn't to say that marketing and sales don't also spend a considerable amount of time understanding the market. They do, but the product team contextualizes everything in a way that prioritizes what gets built and physically manifests something to promote and sell through marketing and sales. The ultimate goal of building a product strategy, choosing the optimal metrics and KPIs, and getting alignment from your leadership team and main stakeholders is to make your product successful – said in another way, to ultimately achieve **product market fit**. How does product-led growth relate to this?

The best combination of factors is for the product function to take in feedback from the marketing and sales functions in an organization (among the many other sources of feedback), digest that information, and apply it to what's being built. Then, once it's gone through that wave, the product function sends out new communication and updates on how the messaging and the product can be improved and optimized. Ideally, this process will happen based on the product experience itself. The importance of fostering an environment that supports the pursuit of product-led growth within the business can't be understated for AI/ML PMs.

In a 2018 study by Epsilon, we can see that 80% of consumers are more likely to stick with a brand when it offers them a personalized experience. A more recent McKinsey study shows that 71% of consumers are starting to expect personalization, and 77% get frustrated when it doesn't happen. The companies that do personalization well aren't just delighting customer expectations, they're growing 40% faster because of it. According to Forrester, "*68% of B2B buyers prefer doing business online versus with a salesperson, and when they engage with sales, they want that experience to be in a more problem-solving, consultative manner.*"

The trend toward an increasingly digital market is resulting in a fundamental change to the way we build and ship products because the way we build products must change to meet this trend in both B2B and B2C domains. Customers want experiences to be personalized and for offers, products, and services to be accessible to them, but when they do interact with a person, they want that experience to be meaningful. It's not a matter of internal politics. It's a matter of responding to macroeconomic trends that impact all businesses across the board. Products must be built with a high level of technological optimization, whether they're AI/ML products or not.

The tech stack we're going to explore in the following section isn't even necessarily a matter of preference; it's increasingly becoming a requirement of product teams. We can't build and optimize a product if we're feeling around in the dark making decisions on hunches and antiquated business acumen. We're not proponents of fully automated decision-making. We're passionate proponents of a collaborative approach in which PMs, technologists, and leaders use their critical thinking and long-term planning skills to make decisions that are backed by data.

It's also not a matter of unnecessarily surveilling users and customers, but of consent. If they want personalization, they can't experience that without a large degree of data collection, analysis, and recommendation. If customers and users want a curated, guided experience in which they can see the value and relevance of a product as soon as it arrives on the scene, product teams will need to become intimately aware of what's happening under the hood. As long as customers and users are made aware that their own data is powering functionality, trust and transparency can be built. At the end of the day, the product team largely takes ownership of this data because it's directly relevant to the work they need to do, which is building and shipping products.

The cornerstone of product-led growth centers on the idea of **value**. If you've built a product that's instantly able to communicate to its user that it will make their lives much easier and they know that because they can start using it right away, you've understood the assignment. How you get there is a *choose-your-own-adventure* story. Sifting through the discombobulated mud that is a company's data heap will involve a rich assortment of platforms and tools. How many products you end up adopting to help you on this adventure will depend on how committed you are to your product-led growth-hacking journey.

Now, let's take a closer look at what kinds of tools PMs rely on to steer their ships.

The tech stack – early signals

Understanding whether your product works for your customers and users or not can be difficult or delayed when you are not set up to get an early signal. Not all companies have the luxury of asking their customers and users directly for feedback that's comprehensive enough to use to inform decisions. This is why investing in a growth-hacking tech stack is helpful if you're trying to understand whether your product resonates with your audience, particularly if it's an AI/ML product.

If your organization isn't set up to make the most of the data you do have, you'll have to rely on ad hoc channels. These are likely to be direct feedback from customers or a noticeable absence in sales, both of which are unreliable measures because they may not reliably give you true insight into the *why*. Think of the categories of tools we will discuss in the following subsections as foundational building blocks you can put in place to set your product and company up for success, whether your product is built for other businesses (B2B) or direct consumers (B2C).

We can't emphasize this enough: investing in the tech infrastructure that will give you the early signal you need to build confidently will save you time and money in the long run. You don't have to involve everything in the following list, but if you want to set yourself up sustainably, it's recommended.

Customer data platforms (CDPs)

CDPs are platforms that collect and centralize your customer data from multiple sources into one view for each of your customers so that you can understand who they are, how they find you, what drives their behavior when they do interact with your product, their transactional data, and demographic data. They allow you to leverage the internal data you have into customer profiles so that you can better explore the kinds of customers you're serving with proof. You might already have channels in place to maintain some of your customer data with systems such as Salesforce and HubSpot, but generally, CDPs also centralize data from those sources. Their purpose is to collect, unify, and manage your customer data so that you can act on it.

Marketing, sales, and your go-to-market team will be important partners for you as an AI/ML PM, so building strong relationships between these teams through the use of a CDP will be an important step toward making sure your customer base trends in the right direction. If it doesn't, CDPs will allow you to collaborate with your business development teams on new campaigns to make sure the value of your product is adequately communicated.

Again, because of the high operating costs associated with AI/ML products, you'll want to develop strong relationships with these externally oriented teams. A big reason for this is to make sure your product–language fit resonates with future and current customers so that you initially and continually see value coming from the product you've built. This is an important part of establishing an AI/ML-native product because these products are newly entering the market. You won't yet have an established baseline of what works and what resonates most with your potential/current users. Expect to spend some time on trial and error with your sales, marketing, and go-to-market teams to get that right so that you can carry on and build with confidence once you have a secure enough foundation.

Reputable CDPs you can explore include Segment, Klaviyo, Hightouch, Insider, and Census. The following are a few considerations to keep in mind when choosing a CDP:

- Ensure the CDP can integrate (or has APIs/connectors) with all relevant data sources that are necessary for your product/business, particularly marketing tools like email platforms, ad networks, and social media sites.
- Ensure the CDP works will all types of data, ranging from structured (labeled, relational) and semi-structured data to unstructured data (unlabeled text, images, videos, audio, or behavioral).
- Vet the CDP's ability to create unified customer profiles, empower data quality, and handle duplicates, particularly as you scale your business.
- Review, CDP's segmentation and personalization capabilities through customer behavior cohorts, advanced segmentation, and predictive insights, particularly those that have built-in predictive analytics models and ML models for churn prediction or **customer lifetime value** (**CLV**) prediction.
- Determine which pricing structure is best for you as you're looking at options. Some are based on data volumes while others are based on the number of users or other features.

Customer engagement platforms (CEPs)

CEPs encompass a wide category. They are essentially any platforms that allow your customers to reach out to you and for you to reach out to your customers. They allow you to do everything, whether onboarding your new customers or users when they first start working with your product, sending personalized in-app messages or welcome messages, setting up interactive walk-throughs for your users and customers so that they get to know your product, getting them to achieve certain behaviors or milestones with your product, announcing new features or new use cases to your customers and users, or sending out links and driving traffic to certain parts of your overall product experience. Here are the key benefits of these platforms:

- Whether you are a B2B SaaS or B2C company, you'll want to create some way of connecting with your customers directly beyond an email list. Let's face it, most of us don't check those emails as often as companies would like. They lack relevance, and the chances are if a customer or user is already in the environment of your product, either in the app or platform, they want to be there and they want to know relevant information about their journey. As a PM, you should concern yourself with the **user experience** (**UX**), user journey, and adoption of your product. CEPs are a great way to create intimacy with your customers and guide them as well.
- CEPs are also a great way to reduce how much customers reach out to your support teams because they're confused about how your product works or what they can do with it.
- They also help drive revenue growth because if there are additional opportunities to upgrade or sell your customers premium features, CEPs help make those options more visible to your customers.
- As an AI PM, you should be aware of the features that will most delight your customers, so using CEPs is an effective way of managing your customers' behaviors in-app in a way that gives them the best experience possible. They're also a great way of gathering feedback through in-app surveys you can deploy throughout your product.

Reputable CEPs include UserPilot, AppCues, WalkMe, and Intercom. Here are a few considerations to keep in mind when choosing a CEP:

- Ensure the platform can support all relevant data sources, CRM systems, marketing automation tools, and channels like email, social media, and SMS, particularly if you're looking for real-time offers and engagement.
- Like CDPs, ensure the platform can give you unified views of customer profiles that are subject to segmentation, targeting activities, and personalized content delivery.
- Vet the platform's ability to help you with A/B testing various campaigns and to provide insights relating to customer behavior and engagement metrics.
- Understand the level of support from the community the company is nurturing to understand if you'll have dedicated account managers, technical support, or onboarding assistance as some of these platforms can be complicated to make the most use of.

Product analytics tools

Communicating with customers and hearing back from them is a huge help in maintaining an open dialogue and keeping the feedback loop alive so that you can continue to build and push out features that customers want, as well as improve features that are hard for customers to interact with and see the value of. However, suppose you don't want to tell your customers where to go. Instead, you want *them* to show you where they go, what they do, and how long they do it for. This is where product analytics come in handy.

Product analytics tools allow you to track all sorts of behavior and create funnels of users to better understand how they physically navigate through your product experience. You can track your customers and users when there are certain events in your app or platform experience, analyze those events en masse, put triggers in place to alert you when certain users perform certain actions, and segment those users so that you can streamline and personalize your communications to them further and create funnels of their journeys.

Typically, product analytics tools come with some built-in KPI tracking and dashboards and let you visualize certain flows and funnels as well. However, if you have a data warehouse in place, you can also send that data over to your warehouse and use a BI tool to analyze it further or pair it with other datasets you might already have in your warehouse.

A typical AI/ML PM is likely to work regularly with these tools because this is where you can keep your finger on the pulse of your product and get enough of an early signal to at least confirm whether your customers intuitively interact with your product and move through your systems in the way they were originally designed. Even if you don't use a CEP, a product analytics tool can still give you tons of insights into whether or not there are glaring issues from the outset. As the product continues to mature and evolve, it should also show you whether or not your customer and user base is evolving along with it. Reputable product analytics tools include Amplitude, Pendo, Mixpanel, Matomo, and Heap. Let's discuss a few considerations to keep in mind when choosing a product analytics tool:

- Get a full breakdown of what the feature set is because many of them will vary in terms of what they can monitor. Things like event tracking, user journey analysis, cohort analysis, segmentation, filtering, and funnel analysis are all features that will help you understand how users are navigating through your product experience.
- As with CDPs and CEPs, discussed earlier, make sure all your data sources can integrate with the tool you want to build. In many cases, they will have APIs and connectors to help make those connections if the product falls short.
- Many of the insights will be used outside of the product analytics platforms, so make sure data can be imported and exported for further analysis.
- Onboarding and implementation can be robust for certain product analytics tools. Things like setting up tracking, creating events, and configuring dashboards can range in complexity from product to product, so make sure you have time to look into that before choosing.
- Ask about specific data volume-related performance issues; certain tools can have challenges with large amounts of data.

- Make sure dashboards can handle your most relevant product KPIs and metrics and that custom reports can help you dive deep into certain aspects of user behavior and product performance.

A/B testing tools

We've discussed the concept of optimization at length throughout these chapters, as well as the spirit of experimentation that's optimal for a growing AI/ML-powered product organization. This is perhaps best captured with A/B testing, a concept we first introduced in *Chapter 1*. While there's a fair amount of comparison when it comes to the types of models you might employ when building your product with AI/ML features, you will also need to test certain versions of your product with different groups. You'll likely be A/B testing all sorts of things – how your product looks and feels, which features you put where, how you guide your users, how you track metrics, and which metrics you actually track – to inspire the kind of product adoption that will have your audience evangelizing about it on your behalf.

It can be hard to make concrete decisions about what will appeal most to your user base when you design a product, particularly when it's a novel AI/ML product that's being released to the market. A/B testing allows you to collect data on which iterations are more successful than others. There are so many routes you can take to optimize the use of your product, how personalized it is, and eventually decide on which iterations best lead to conversions and further revenue. Investing in A/B testing tools not only reinforces a culture of experimentation but also helps you track countless tests and stores the insights from those tests for you.

You can once again choose to send this data out to your warehouse and further mine it for insights and append it with other methods of experimentation from the other tools in your BI tool, or you can store it on the dashboard of your tool of choice. Either way, A/B tools allow you to organize and keep track of the results of your efforts, allowing you to do hundreds of A/B tests if you need to. You can use these tools to A/B-test anything, whether marketing campaigns, features, buttons, links, colors, fonts, or anything that you can derive quantitative and qualitative feedback from. Reputable A/B testing tools include VWO, Optimizely, AB Tasty, Google Optimize, and Five Second Test. Let's take a look at a few considerations to keep in mind when choosing an A/B testing product:

- Ensure the tool can support various experimentation types to accommodate different testing needs.
- Consider options for segmenting users based on behavior, demographics, segments, and other criteria to understand the extent of tailoring capabilities.
- Check integrations with personalization activities you might be doing elsewhere to see if you can combine efforts within the platform.
- Gauge how much technical expertise is needed for implementation, onboarding, and setting up your first experiment and whether you have the in-house expertise required to handle it.
- Look at the statistical methods used to analyze the results of your experiments.
- Confirm the sample size minimums for your tests to make sure you can conduct experiments that will yield statistically significant results.
- Confirm the tool's ability to allocate traffic for experiments according to your needs during tests.

Data warehouses

A data warehouse is a kind of relational database that allows you to centralize data much as a CDP does, but the main advantage here is that it formats, transforms, and standardizes that data so that it can be made available for things such as the CDPs we discussed earlier, as well as BI tools, which we will look at next. You might recall that we covered the importance of data warehouses in *Chapter 1*. It suffices to say that you'll probably need a data warehouse because it serves as the backbone of all relevant company data, customer or otherwise. Think of it as the main artery that pushes data out to various tools. It will also allow you to query the data you have so that you can gain insights from it.

Your data warehouse will be an agnostic place in which your data will live when it's not actively being called through a direct query or through one of the other tools you hook it up to. When you manage a native AI/ML product, you will think about the success of your product agnostically as well because, the truth is, success comes from a collection of events. Your product needs to be communicated about with customers and users, but it also needs to be built in a way that directly solves their problems. Your data warehouse will serve as your main source of truth, sending data appropriately to the teams that deal with communication, as well as to the teams that deal with building.

Keep in mind that how you store your data is going to be a pretty big financial and strategic decision. You might opt for a data lake or other relational/non-relational databases that aren't considered to be data warehouses – folks in other roles (like solution architects, data engineers, and consultants) will probably be making that call. As an AI PM, we would advise you to be a foundational part of that conversation because how and where you store your data will have downstream impacts on things such as which other growth-hacking tools you end up adopting and ultimately how much insight you'll be able to gather from your data.

If your data is in a data lake, you probably won't interact with it very much. Data warehouses are built for the express purpose of refining your data and getting it ready to push out to other contexts easily. Reputable data warehouses to consider would be Snowflake, Amazon Redshift, Google BigQuery, and Databricks. Below are a few considerations to keep in mind when choosing a data warehouse:

- Ensure the data warehouse you intend to use will scale with your data volumes, processing needs, and storage needs, particularly if you're setting one up for the first time.
- Have an understanding of how fast you need the query to be, particularly if you're setting it up for advanced analytics and machine learning to inform fast decision-making.
- Look for built-in ETL tools to refine your data and, if they don't exist, confirm they have third-party ETL tool integration because you will need to refine your data to make it usable for the tables in the warehouse.
- Assess data modeling options to make sure the warehouse will support your current data structure, schemas, and querying needs, as well as data types (structured, unstructured, semi-structured, etc).
- Understand the total cost of ownership of the data warehouse, which should include onboarding, setup, storage, compute, and additional services.

- Ensure compatibility with the BI tools you will be using to extract insights from your data.
- Confirm backup and recovery options to prevent data loss in case of incidents, outages, or data breaches.

Business Intelligence (BI) tools

While your CDP and your data warehouse organize data in a certain way for you to consume, they probably aren't going to be robust enough for you to explore it effectively. While CDPs are nice for establishing some patterns and understanding the journey of your personas and so forth, they aren't intended for the exhaustive exploration of your data. You can also technically query data from your data warehouse directly, but you won't visualize or wrangle that data directly in the warehouse. That's where BI tools come in.

BI tools allow you to take all the data that's in your data warehouse, data that's coming from all sorts of places, and actually analyze it. The purpose of a BI tool is to help you answer important business questions about your data. Data itself, lying in a dormant state, is not helpful or insightful. You have to analyze it to get the treasure, and in this case, that treasure comes in the form of trends, insights, and knowledge. Taking data from a dormant, disparate state, unifying it, and processing it is a tedious process that's very hard to do without a BI tool. BI tools also allow you to create dashboards for teams to use so that they routinely get a health check on the metrics we outlined earlier in the chapter.

While your data warehouse might serve as your source of truth, you will use your BI tools to understand whether that truth actually reflects reality because it's virtually the only place you can actually see it. With a BI tool, you don't just check the validity and visual expression of your data – you create a system of reliance on making data-driven decisions. There will always be an element of risk in the creation and expression of a new product or business. You don't know how it's going to go. BI tools help us see what we do know as clearly as we can and help us minimize that risk.

BI tools aren't just super valuable internally. They also allow us to make visual representations of our data that we can communicate in the form of customer dashboards or through marketing and sales collateral about our products. If you create infographics and sales decks and drive customer communication that involves charts and graphs of your data, they are probably all made through the use of a BI tool.

In essence, a BI tool works as a refinery for your data so that you can further craft the message you want to send out, as well as the message you want to reflect on internally. Notable BI tools include Power BI, Tableau, Sisense, ThoughtSpot, and Looker. The following are a few considerations to keep in mind when choosing a BI tool:

- Assess how user-friendly your BI tool needs to be. Does it need to be accessible for people with low technical skill levels, or can it be used by more technical members of the team?
- Understand your dashboard, visualization, and reporting customization needs, particularly if you have niche or specific needs for your dashboards. Not all BI tools uniformly allow you to showcase your data. Certain ones are better than others when it comes to how you show data. Take stock of different needs at the department level because some departments may have different needs than others, but the tool needs to work for everyone.

- Check to see what level of data preparation and cleansing the BI tool offers. This will integrate with your data warehouse, and if neither tool can support your specific data ETL needs, you'll need to handle that independently.
- If interactive dashboards and embeddings are important to you, ask about how the tool will handle that because these are not default features. If you need to drill down into data for deeper insights, you may not be able to with certain tools. If these features are available but for an additional cost, understand those costs fully.
- BI tools allow data to be shared and exported for various purposes, so understand what options you have for sharing reports and dashboards with team members and external stakeholders.
- Vet the level of advanced analytics or NLP capabilities your BI tool will offer to query your data in the platform, particularly for users that are less technical.
- Look for BI tools that have robust communities around them for things like user guides, tutorials, and customization tips and tricks from power users. In many cases, BI tools do not have a comprehensive, dedicated headcount to help you set up these tools if you don't have the technical capabilities in-house.

Growth-hacking tools

Finding quick success isn't just the imperative of an AI/ML PM but of all PMs and entrepreneurs. Figuring out the best way to make money, increase brand awareness, and find quality leads is what growth hacking is all about. Although most products won't go viral overnight, there is some method in the madness of investing in the tools we went over in previous sections. They all get you closer to the information you need to make present and future decisions in a way that's informed by what works and what your customer truly wants from you.

You can imagine growth hacking as part of your go-to-market strategy and team and, often, it's discussed as a set of tactics that help scale and grow the adoption of your product. The idea is you try a set of practices and techniques to see if they have an early impact on your metrics and adoption. You often try a selection of strategies to see which is the most impactful, and this includes quite a large degree of experimentation. This allows your marketing and business development teams to get creative, strategic, and agile with their approaches to their own work and gets them thinking outside of the box.

Growth hacking can mean a lot of things. Sometimes, it may involve using product analytics to measure the success of certain strategies against performance or engagement. Other times, it could mean testing different marketing channels, messages, or social media tactics to see what line of messaging lands best with the customer and user base you're targeting.

We went over some broad categories of products that certainly help with growth hacking and finding success with your AI/ML product previously, but there are so many valuable tools out there that don't fall into those categories per se:

- Products such as Expandi allow you to use LinkedIn for social selling campaigns.
- Crystal Knows uses AI to craft personality profiles to provide insights into behavior and sentiments.

- Landbot helps you build chatbots to interact with your customers in-app or on your platform.
- Hotjar allows you to see heat maps and other analytics of where users' cursors move in your product.
- Lyssna helps you conduct UX research with real users.
- Fomo helps you build credibility with your brand through transparency and social proofing.
- Leadfeeder helps you turn your page visitors into leads.

There are endless products out there to help companies achieve the right balance of product market fit. The route you choose will be a combination that keeps the gears of development going but that also allows you to stop and reflect. Processing choices that have been made allows you to assess whether or not those choices are bringing you closer to your ultimate goal of successfully commercializing your AI/ML-native product.

In the following section, we will discuss the costs that contribute to the expense of managing AI pipelines, as well as their impact on pricing. While incorporating AI might improve efficiency and should, theoretically, make your price point lower, the total cost of managing an AI program internally is quite high. That cost is often passed onto your customers because it contributes to the overall cost of running your product. A few considerations to keep in mind when choosing a growth hacking tool are:

- Understand how complex and time-intensive the onboarding process is, particularly for using features without extensive training.
- Understand how compatible the tools are with your existing tech stack for seamless data flow (and API availability).
- Assess the built-in testing features to experiment with different growth strategies and measure their effectiveness.
- Check to see what kind of user segmentation capabilities it has when you're testing different variations on specific groups.
- Confirm the level of team collaboration the tool will enable for you to share access to campaigns, reports, results, and insights from various tests. Does it empower members of various teams to comment and give feedback?
- Check the level of automation the tool offers for things like automated drip campaigns to engage users over time.
- Assess version control capabilities that allow tracking changes in campaigns and access to previous iterations that worked well.

Managing costs and pricing – AI is expensive

As you may recall, we discussed the costs associated with building an AI/ML program in *Chapter 2* and *Chapter 3*. Formulating a pricing strategy will be a highly personalized experience that will involve a number of factors, from the comparative prices of your competition to the operating costs for managing your AI/ML infrastructure and workflows. In this section, we will briefly cover the various aspects of AI product management that impact costs and how to use this knowledge to inform your pricing strategy so that you're aware of the main contributors to your AI/ML program costs.

Given the high cost of operating, you'll want to be sure that your product serves the needs of your market and your customers as quickly as you can. Keep in mind that your costs are likely high when managing an AI/ML program, but they're also dependent on how many AI features you've built into your product, particularly if the underlying logic powering your product as a whole is driven by AI/ML. The longer you operate in the dark, the more resources you'll spend on your product. This means you'll know later rather than sooner whether it resonates with your customers.

Let's first start with the cost of AI/ML resources themselves. According to WebFX, most AI consultants charge between $200-$350 an hour and the cost of a custom AI solution is anywhere from $6,000 to $300,000. Using third-party AI software instead can cost anywhere from $0 to $40,000 annually. Real costs are subject to change based on the availability of talent, reputation, and location, but these are reliable ballparks for now.

It might be tempting to build an AI/ML-native product around consultants, but it's not the best practice. As soon as something goes wrong, you'll spend copious amounts of money fixing even modest problems. If you want to hire someone in-house and avoid having to use outside help, you're looking at an average salary of $160,000+ for an ML engineer in New York City. If you want someone to manage a team of ML engineers, you're looking at an average salary of $200,000+ in New York City.

Then, there's the AI enablement tech stack. Running, maintaining, training, querying, storing, and processing data and AI systems all have costs associated as well. That doesn't include many of the growth-hacking tools we have already mentioned, which all vary in price as well. Managing all the various vendor relationships, contracts, pricing bands, and caps so that your usage doesn't grow astronomically will probably warrant its own head count on top of it all. Managing AI/ML infrastructure and empowering a product team to scale and grow to meet the market that it services is an expensive endeavor, but that doesn't mean it's not worth the pursuit!

> We're not going into too many specifics here since your AI strategy will largely determine which tools, platforms, partners, and roles you hire for. You'll collaborate with key stakeholders to decide on suitable options based on your specific use case and budget.

Case study

Let's now return once again to our Waterbear case study to see a practical example of how benchmarking, performance, growth hacking, and costs are balanced for the production of their flagship product, Akeira. Earlier in this chapter, we looked at a list of KPIs and value metrics. In this section, we will be expanding on some of these metrics and discussing how these metrics were vetted, selected, and prioritized. We will also go into the OKRs that drove the organization to manage the performance of Akeira.

North star metrics

Akeira's product management team considered a number of metrics to include as a north star metric. These metrics included:

- Engagement metrics like **daily** and **monthly active users** (**DAU** and **MAU**)
- Interaction metrics like frequency of use and number of entries per user
- Quality metrics like user satisfaction and recommendation relevance
- Impact metrics like self-reported outcomes

But, in the end, they chose a metric that strongly aligned with the mission and values of the company for their north star metric: the **monthly mental health improvement index** (**MMHII**). This also aligned with their product strategy, which was to allow the experience of Akeira to be product led. For Akeira, that meant gathering in-app feedback based on how the product was functionally improving the lives of its users.

Waterbear's founders wanted to choose a north star metric they could base their product decisions on, and they set out to create the company for the express purpose of improving the lives of the users. They demonstrated this when choosing their north star metric, which revolved around a personal MMHII as well as a collective MMHII.

This choice was also strategic in the sense that it separated the in-app experience from outcomes and well-being. All tech companies want to build a world-class app experience based on user preferences and behavior that keeps users coming back. But what does it mean to have an amazing app experience if it doesn't functionally improve the lives of its users? Choosing MMHII meant that they were prioritizing proven outcomes above the in-app experience. Doing so also meant that they could align their internal decision-making and external communications around proven mental health outcomes, which resonated highly with their customer and user base. They were speaking directly to their community by doing this.

KPIs

Other KPIs that were heavily tracked were user engagement, user interaction, quality, impact, and retention metrics. These metrics were able to split the product experience into various areas to assess performance and success, and to help provide early signals when the product was in danger of suffering in one of these areas. Each of these areas also adheres to the positioning of the north star metric. If users are not engaged with the product and if they're not interacting with it regularly, they won't see a positive outcome in their mental health given the nature of feedback loops that impact our mental health positively.

If the quality of the interactions they're having in the app aren't high, their engagement will suffer. If the impact of their mental health improvements can't be quantified and defined in various areas, we can't claim to be impacting their MMHII positively in a meaningful way. Finally, if users can't be retained, that also means their mental health will not be served. The operative word here is "served." If your KPIs are not in service to your north star metric, they will only represent superficial successes and should be considered vanity metrics. Vanity metrics might fool an immature investor and build short-term success, but they're not a strategy for sustainable growth.

Here is an overview of the established Akeira KPIs we saw in this chapter, from a product perspective:

- **User engagement metrics:** Engagement metrics allowed the product team to prove Akeira was product led because they established a pattern of behavior of users returning to the platform for sessions. A steady increase indicated onboarding was effective and that users were finding consistent value in the app. Longer sessions showed that users were finding the content and features engaging. Frequency of use was helpful for the Akeira product team to segment various types of users and to work on more personalized experiences for users of different ages and backgrounds. Here are some metrics in this category:
    - **DAU/MAU:** The number of unique users who journaled with Akeira on a daily/monthly basis
    - **Session duration:** The average time each unique user spent on the app for each session
    - **Frequency of use:** How often they came back to the app over time (daily, weekly, monthly)
- **User interaction metrics**: User interaction metrics helped the product team demonstrate product effectiveness. They cross-referenced these users with those who submitted support tickets to confirm that the app intuitively guides most users, retains users, and promotes organic growth demonstrated by the retention rate. Tracking the number of entries helped confirm the app was helping users develop a journaling habit that served the north star metric, and the average length of entries demonstrated the depth with which users were connecting with Akeira. Here are the key metrics in this category:
    - **Number of entries:** The number of journal entries created by each Akeira user
    - **Average length of entries:** Average word and character count for each journal entry for each Akeira user
    - **Retention rate:** The percentage of users who continued to use Akeira over time (daily, weekly, monthly)
    - **Support ticket frequency**: The number of support tickets opened by active users
- **Quality of interaction metrics:** Quality of interaction metrics centered around proving the efficacy of the ML models that were powering Akeira, most notably NLP models and the recommendation engine that powers suggestions. These metrics also helped confirm user trust in Akeira's AI capabilities. High user satisfaction metrics meant that the app had referral potential, increasing the likelihood that users would recommend Akeira to friends, which also furthers organic growth. Key metrics in this category include:
    - **User satisfaction:** Built off an in-app survey to measure user satisfaction and how likely each user was to recommend Akeira to their friends.

    - **NLP analysis accuracy:** These metrics were built on the reward systems in the NLP that underlie Akeira. They measured how accurate Akeira was at sentiment analysis, topic extraction, and entity recognition for each user.
    - **Relevance of suggestions:** These metrics graded the performance of Akeira's suggestions or prompts to users as they were reflecting.

- **Impact metrics:** Impact metrics are aptly named because they made the greatest impact on the MMHII, which helps Akeira deliver on its promise to users. This is because demonstrating user transformation was essential to proving meaningful improvements. Self-reported outcomes also added qualitative data to add depth to quantitative metrics like the MMHII, much of which was organized and analyzed for building compelling marketing stories and messages. Let's look at the important impact metrics:
    - **Established behavioral changes:** These metrics captured changes in the users' behavior, attitudes, or emotions over time based on the analysis in the journal entries. They were reinforced and confirmed by the users.
    - **Self-reported outcomes:** During the confirmation process for the above metrics, users were asked for open-ended feedback for further context on the perceived impact of journaling on their well-being, self-awareness, and personal growth.
    - **Goal achievement:** Though Akeira tried to recommend and notice goals on its own, these goals were confirmed by the user. Once they were confirmed, Akeira worked to substantiate progress toward those goals based on the journal prompts and, once the goals were achieved, Akeira celebrated those achievements with its users.
- **Retention metrics:** Retention metrics helped the Akeira product team understand which features or challenges customers had that led to them leaving the platform so that they could focus on minimizing those challenges for other customers. Direct feedback from customers allowed the team to address areas where user expectations might have surpassed app functionality. Some critical retention metrics are:
    - **Churn rate:** The percentage of users that stopped using Akeira within a specific period of time (weeks, months).
    - **Reasons for churn:** To best inform product improvements and strategies for retention, Akeira provided a survey to those who wanted to give feedback to help improve the app experience.

OKRs

Next, let's take a look at Akeira's OKRs:

- **Objective: Improve mental health outcomes**

  This is aligned with Akeira's north star metric, which is again a reflection of Waterbear aligning its product strategy with the goals it has established to gauge product success. By focusing its top OKR on MMHII, and reflecting that mental health improvement against the self-reported mood, anxiety, depression, and sleep quality, they put the focus on the purpose for which the app was created. This is also an area Waterbear focuses its marketing efforts on, whether it's blog posts, deep dives, fireside chats, or white papers. By defining mental health outcomes and what success looks like and sharing it, they're ensuring their product strategy leads with their values. Here are the key results associated with this objective:

  - **KR1**: Increase average MMHII per user by 30% within a year.
  - **KR2**: Increase the percentage of users reporting mood improvement by 25% within the next six months.
  - **KR3**: Achieve a 15% reduction in self-reported anxiety and depression symptoms among users within a year.
  - **KR4**: Increase the percentage of users reporting improved sleep quality by 30% within the next nine months.

- **Objective: Enhance conversational effectiveness**

  This OKR focuses on the tech powering the app experience. Here, the completion rate refers to the interactions by users that are successfully "finished" without the user abandoning the discussion because they don't feel it's worthwhile. Conversational AI apps can be finicky to manage and, even when they are effective, they struggle with retaining user interactions. By placing conversational effectiveness at the second level of priority, Waterbear is sending a powerful message to users and internal stakeholders alike: we don't deserve to maintain a product if it's not effective and if the underlying tech that powers it is not improving over time. The following are the key results expected:

  - **KR1**: Improve completion rate to 70% within the next two quarters.
  - **KR2**: Maintain a user satisfaction score of 4 out of 5 or higher consistently.
  - **KR3**: Reduce issue resolution time to under 1 minute on average by the end of the year.

- **Objective: Increase user engagement**

  Finally, the third OKR is exemplary of what you might see with many traditional software apps: ensuring users are engaged and keep coming back. By making this the third OKR, the company is signaling to everyone internally and externally that app engagement is tertiary to proven mental health outcomes and enhancing tech effectiveness. Often, we see tech companies cannibalizing their own success by prioritizing in-app user engagement before the outcomes for which the app was created in the first place. This offers a ripe environment for the establishment of vanity metrics. Given how hard it is to shape company culture even when leadership is aligned, it can be hard for an organization to come back from this culturally. The key results are:

  - **KR1:** Increase DAU by 20% within the next quarter.
  - **KR2:** Increase average session duration by 15% by the end of the year.
  - **KR3:** Achieve a retention rate of 40% for users within the first month of Akeira usage.

Growth hacking

In earlier case study sections, we discussed Waterbear focusing on beta users and testers before they scaled their product. In this section, we'll discuss what that period of time looked like practically, because it was the time when growth hacking was most in focus for Akeira. After testing the product with a very small group of friends, colleagues, and beta testers who signed up for the mission of using tech to address the mental health crisis impacting the globe, Waterbear got to work on testing assumptions and messaging with a wider audience to see what resonated best.

They wanted to drive organic growth, so they started their growth hacking journey by focusing on an incentivized **referral program** offering considerable discounts to friends of friends who joined. Waterbear did considerable A/B testing with various features, charitable donations, and cost reductions for joining to see what resonated best with that user base. By leveraging social proofing, they were able to create a reputation around the app's effectiveness. Because there were not a lot of apps like it in the market, and the apps that did exist were only capable of handling certain competencies of Akeira instead of the full package, the social proofing paid off.

Part of their social proofing strategy was built on **social media engagement**. By encouraging users to share their stories on TikTok, Reddit, and Instagram, particularly early on, they were able to create a community around their app's effectiveness and success. This built a lot of trust in their small but mighty community, and growth came from organic means. Building on the social impact, they also created a trend on social media where users were sharing what they love most about Akeira, personifying the app experience further. This gave Akeira a brand identity and personality with current and future users of Akeira, which became the mascot for Waterbear in its own right.

This community influence blossomed further through the use of partnerships with influencers who had also built a personal brand around mental health and the importance of using tech to serve our collective human growth and not the other way around, which also created a splash. This was unexpected for the PM team for Akeira, but it started a wave of influencers and social media users criticizing tech companies that were not serving the markets they claimed to be. Nurturing a community based on **authenticity** and **transparency** was by far the greatest tool in Waterbear's growth hacking shed. It also allowed users to discover new ways to leverage Akeira and share that feedback with its PM team.

Remember that growth hacking isn't just about acquiring more users or even making more money. It's about creating a feedback loop of sustainability and growth that helps a company scale. In Akeira's case, the growth hacking did lead to higher user engagement and sales. But it also led to an additional well of inspiration: valuable user feedback that improved the app experience tenfold.

Summary

In this chapter, we focused on the tracking, marketing, promotion, and selling of the AI/ML product. We covered the various ways that an AI PM can benchmark and track their product and its success using metrics and KPIs, as well as what that means for the greater organization and the successful adoption of that product among its user base. We also contextualized this benchmarking against the overarching product strategy and vision that powers what gets tracked and measured. All these activities help PMs get internal signals into whether or not the product they've built works for their active customers and users.

Then, we discussed the greater work of getting external signals for what is and isn't working using the various tools in the growth tech stack that directly connect to the UX. We went over the elements of growth hacking. Whether you're looking to optimize how you acquire, engage, or retain customers, you're going to have to find a way to gather that data, analyze it, and use it to make real decisions about how to add, remove, or improve features. Overarching, long-lasting trends of organizations becoming increasingly data-oriented and data-driven have been massively influential drivers of creating a new culture in the software world.

Customers increasingly look for an experience that doesn't just feel intuitive and natural, but that reinforces to them over and over again that they've come to the right place and that the product they're interacting with is the right tool for them. They don't want to have to speak to a live person and they don't necessarily even want to take what marketing is selling them. They want to interact with a product directly and see whether its inherent value is apparent through their own exploration. This is a massive shift from how products were built and sold in the past, and this means that our responsibility as AI PMs is to offer an experience that aligns with how customers want to be regarded.

We're delighted that we've entered a new era in which we aren't playing coy anymore. Customers and users know they are being tracked and they're putting us to the test. Customers and users want us to use our intelligence with them in order to give them an experience that won't insult their intelligence, and we're here for it. It's a brave new *help me, help you* world, and we think it's about time we use all this data to make the UX simpler than it has ever been.

In the next chapter, we will cover key considerations of managing AI-native products. We will discuss managing alignment through leadership and communications, managing people through creating an empowered, safe atmosphere to learn and experiment, and all the key considerations when managing operations.

References

- *Landbot*: `https://landbot.io/`
- *ProductPlan: Product strategy*: `https://www.productplan.com/glossary/product-strategy/`
- *Choosing your north star metric*: `https://future.com/north-star-metrics/`
- *ML Manager salaries (Salary.com)*: `https://www.salary.com/research/salary/hiring/machine-learning-manager-salary/new-york-ny`
- *ML Engineer salaries (Built In NYC)*: `https://www.builtinnyc.com/salaries/dev-engineer/machine-learning-engineer/new-york`
- *AI pricing*: How much does artificial intelligence cost: `https://www.webfx.com/martech/pricing/ai/`
- *Leadfeeder*: `https://www.leadfeeder.com/`
- *Fomo by Relay*: `https://fomo.com/`
- *Lyssna*: `https://www.lyssna.com`
- *Hotjar by Contentsquare*: `https://www.hotjar.com/`
- *Crystal*: `https://www.crystalknows.com/`
- *Expandi*: `https://expandi.io/`
- *The death of a (B2B) salesman*: `https://www.forrester.com/what-it-means/ep12-death-b2b-salesman/`
- *New Epsilon research indicates 80% of consumers are more likely to make a purchase when brands offer personalized experiences*: `https://www.epsilon.com/us/about-us/pressroom/new-epsilon-research-indicates-80-of-consumers-are-more-likely-to-make-a-purchase-when-brands-offer-personalized-experiences`
- *The value of getting personalization right—or wrong—is multiplying*: `https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personal%ADization-right-or-wrong-is-multiplying`
- *The 'North Star Metric Checklist' and the most common pitfalls*: `https://www.growth-academy.com/north-star-metric-checklist-and-common-pitfalls`

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

11

Managing the AI-Native Product

In *Chapter 6*, we discussed the different stages of the AI-native product development cycle, the dream team that will power your AI-native product, and the way you'd position and sell that product. We also covered the establishment of a product vision, strategy, and roadmap in *Chapter 8*. In this chapter, we'll be expanding on some of those concepts to better understand what the management life cycle looks like after you've launched a product. Managing the AI-native product is a discipline in and of itself and it's what makes the work of a PM challenging. It's one thing to build and enter a market from scratch and quite another to maintain and manage the product thereafter.

As is often the case, hiring managers and recruiters breathe a sigh of relief when a PM comes from an engineering background because they feel it means they will better empathize with the dev team that will support their product. While that's nice, and certainly helpful for more technical PM roles, that's only a small part of the story. When building a team, particularly as an AI PM, I urge you to find candidates who can show their strengths beyond the technical side.

Understanding the strategic elements that power your product choices, from leadership to internal department feedback to the invisible hand of the market, is more important than your technical prowess. Building trust with your colleagues, the people you will depend on to execute the roadmap, strategy, and backlog you build with, is also more important than your technical prowess. Chances are, if you're managing an AI product, your technical understanding of AI concepts is already quite strong. You will need to understand various models and how they work if you're going to work with a team of ML engineers. But as you go on the journey of managing your AI product, you will continuously find that these skills are secondary, not primary.

Some people might be surprised to read that. They may think that as a PM, so much of their work is involved with problem-solving technical issues, and understanding technical debt or the scope of a body of work. This is true; a lot of your work will be focused on this. But if your PM work is largely devoted to these areas at the negligence of the others we'll cover, you're effectively relegating yourself to a DevOps role rather than a true PM role. The core skills for a PM are communication, market research, user research, strategic thinking, analytical, leadership, customer engagement, and collaboration.

We will explore how these skills manifest themselves in the various facets of a PM's work in the following sections:

- The head – Managing alignment
- The heart – Managing people and values
- The guts – Managing the rest

The head – Managing alignment

Defining a product vision, strategy, and roadmap is all part of the work of creating alignment. In *Chapter 8*, we discussed the concept of marrying your product vision and strategy with those of your leadership team. This is because the alignment you establish needs to be a reflection of those of your company. Product alignment is important for the AI native product because AI technology is central to value delivery. AI products need to solve the right problems for customers and they are highly cross-functional. This means all teams involved with development and strategy need to make sure they are working toward the same customer-centric goals.

Failure to do so would render the resource-intensive process of developing an AI system inefficient, which means a lot of wasted time, money, and brain power. We don't work in a vacuum. There's a reason why companies have core values in their career section when they're recruiting for new talent. What's publicly facing needs to reflect the interests, values, and convictions of the greater structure. There has to be a way to set yourself apart from the rest, as well as discerning newcomers in the group. In this section, we will expand on some of the principles we've covered in *Chapter 8* and see how they're employed in practice.

Vision

All organizations amplify certain values internally and externally whether they're aware of it or not. If you know how to look, these values are implicit in a company's brand, messaging, and communications of all kinds. Some companies are better at crafting these communications in a way that harmonizes with their vision and values than others. But this process is happening in the background whether they're good at it or not. This is still the case once you've joined a company and are in the flow of managing a product. If anything, this becomes more prominent when you're an AI PM for a company because one of your primary tasks is bringing clarity and reassurance to a process that's rife with inherent complexity, ethical considerations, and long-term focus.

Reporting structures and role definitions vary at many companies, but irrespective of how your company perceives (and pays) you, an AI PM is always a visionary. You have to be! Visionaries see potential and opportunity where others don't and AI products come with a heavy burden for visionary thought because you're often navigating emerging technologies, intricate environments, and frequently uncharted user experiences. Because these products are AI native, it can be hard to predict how users will engage with AI/ML. What new pain points or benefits of using your nascent product will emerge? How might user needs evolve as your product, or the technology it employs, matures?

We have to be brave enough to consider these potential, imaginary future states and plan accordingly. In order to do that, we must have a vision of what our products' present and future value might be and how to share that vision with others. Through an AI PM's work and strategy, they inspire others to slowly come to see that vision as well. Visionaries that are most effective are able to not only get others to see that vision but also create a plan to turn that vision into reality – where everyone else can see it come to life undeniably. Product visions should be aspirational in nature and they should give people something to believe in beyond the success of the product, or even the company. All PMs want their products to be successful. All companies want to succeed. The question really is: why should yours? What's so special about you?

Beneath story points, product requirement documents, meetings, strategy sessions, and sprints lies a solid bedrock of visionary work. If you're going to be effective at your company, particularly if you're a sole PM or the head of your department, you'll have to get comfortable with being seen in this way.

Good vision statements

Pinterest's vision statement is "to bring everyone the inspiration to create a life they love." This statement is good for several reasons. It's emotionally resonant, inclusive, and empowering, and it puts the focus on a greater global mission, which is focused on inspiration. It calls attention to a cause that's greater than the confines of the product or the company – giving "everyone" a tool they can use to enrich their lives. This empowers users to bring their aspirations to a higher level and gain control over their own lives. The statement also taps into a deep, universal desire for personal fulfillment and creativity, which keeps it rooted in positivity. Because of its inclusivity, Pinterest can make a home for all users on its platform by ensuring accessibility and relevance for people from different walks of life, which encourages engagement and loyalty. There's also a strong, cohesive link between the vision and the product itself, the platform on which users discover and save ideas that inspire them, which builds trust in the platform.

Bad vision statements

Now let's take a less effective example: "We aim to provide the best software in the industry to help businesses succeed." If at first glance this didn't bore you, maybe it seemed vaguely aspirational. And it is. After all, this company wants to be the best! If you've ever taken a walk in New York City, you might have come across countless pizza places that claim to be the best in NYC, or maybe even the world. Surely they can't all be right. "The best" is really in the eye of the beholder and what's the best for one person may be the worst for another. Pinterest wasn't claiming to be the best, just the most compelling. Also, "best software" is quite vague and unclear. We don't know anything about which industry or area of software the company is serving. It also doesn't include anything about what makes the company or product unique. Why does this product or company deserve to be the best in the industry? With this lack of inspiration or clarity, can we have any faith that this business is poised to help any other business succeed? This statement is unclear, uninspiring, and offers no sense of direction or greater purpose. Even the company itself doesn't believe that it will succeed, only that it aims to!

Every word in a vision statement matters because it will be perceived by all employees as well as all potential customers and users. AI PMs know that vision statements can wield a lot of power in shaping external and internal perceptions of your product and it's precisely why you need to play a vital role in helping craft one that sends the message you want. You're the steward of maintaining alignment for your product and there is no alignment without setting a clear vision. People will come to you from all sides with all sorts of product questions. You're the main point of contact for all stakeholders regarding anything to do with your product, so you'll be the one understanding how to contextualize all sorts of requests.

This isn't easy work, especially in the AI space where complexity can compound every day. Entropy is real. With each passing day, week, and month at your company, lots of things can happen to knock you off your focus. The company could onboard a huge blockbuster customer. It could lose a foundational leader or CEO, or experience a merger or mass layoffs. The market could suddenly shift or there could be a sudden shift in technology. The rate at which AI technology evolves only heightens the stakes for AI product teams.

All these factors and more cause an organization and product team to have to rethink what they're doing and whether you need to go back to the drawing board. Sudden shifts in priorities can be costly for AI products and unfocused teams risk wasting valuable resources retraining models, misinterpreting signals, or pivoting prematurely. When technical team priorities become unclear, or worse, reactive, AI product development slows and deteriorates. This loss isn't just functional or pragmatic, it's also emotional. People are not machines, they're our greatest resource. When the very foundation of what you're building, and why, changes, it creates a lot of mental and emotional chaos for the people that are part of your product team and company.

In the face of uncertainty, we have to resist the pull of chaos. We live in a chronically uncertain world, particularly when it comes to AI development and adoption, but we can't be nihilistic about that. If anything, we need to employ tools that create certainty and significance that are personal to us. Defining and communicating a clear product vision is about revolting against entropy and choosing to be rooted in something solid that matters. We aren't leaves in the wind. We're capable, smart, resilient people who want to use the power of our wills to leave a mark on the world. People do it all the time and you can too.

As an AI PM, you possess the tools to create clarity and reassurance through defining a strong product vision that will become the anchor that grounds your team and your stakeholders alike. In an industry that's so defined by change, you can make a commitment to build something that leaves a lasting, significant mark on the world.

The following is a list of questions that can help you create and define a long-term vision for your product:

- What problem is your AI product solving?
- How does it leverage AI capabilities to do so?
- What is the goal of your product?
- What functionality and value does it deliver through ML and automation?
- How do you want your product to evolve over time?

- How will new data or improved algorithms improve your AI product?
- How will your product learn and adapt?
- Who is the target audience for your product?
- How does your product enhance their workflows, decisions, or experiences with ML?
- Does your product address their current pain points and anticipate future needs through predictive analytics or personalized insights?
- How will your users change and evolve?
- How can your AI product evolve with them by adapting to shifts in user behavior or preferences?
- Does your product align and contribute to the overall company mission/vision?
- Does your product support strategic company goals like scaling operations, improving efficiencies, or driving innovation?
- Is your AI product ethical? Moral?
- How does your product respect privacy, avoid bias, or adhere to fairness and transparency standards?
- Are these important to your brand and company mission?
- Does your product foster trust and loyalty among users by maintaining ethical AI practices and ensuring explainability?

Once you have a clear vision for your product, you'll want to quantify whether it's being effectively communicated and executed:

- Can you confidently say that other stakeholders understand the vision for your AI product and how it leverages AI/ML?
- Have you gathered feedback from stakeholders and incorporated their perspectives into your AI product vision?
- Do the metrics, KPIs, and OKRs you've established reflect the success of your product's vision?
- Do you have buy-in from your leadership team for the product vision, particularly in terms of the long-term strategic value AI can deliver?
- Does everyone in your company understand why your product exists in the market, what its competitive edge in terms of AI capabilities is, and how it will enhance user experiences?
- What makes your product special in its approach?
- Are there proprietary algorithms, premium data handling, or specialized AI features that set you apart?
- How does your product compare to competitors in terms of AI performance, accuracy, scalability, or innovation?
- Does your product vision address gaps in your market that aren't served by any competitors, particularly when it comes to how AI is being used to solve problems?
- Does this product vision keep users and their data safe?
- What guardrails have been put in place to protect users and their privacy?
- Are there milestones and deliverables that will guide the product's development that will ensure alignment with the product vision you've established?

We don't build a clear vision for our products because we read an opinion piece somewhere on LinkedIn and think it sounds nice. We do it so that we operate from a position of strength. Executing is not easy. You have to build a plan based on a solid strategy because hiccups come up all the time when you're executing. Building a foundation from a strong vision is a way to prevent those hiccups from throwing you off course so that your work is measurable, ethical, and aligned with both the company's goals as and market demands.

Communication

Reading the previous section might resonate deeply with some. To others, it will sound like abstraction and nonsense. But it's not. AI PMs need to be able to navigate the nebulous gray space of knowns and unknowns, especially in the rapidly evolving field of AI/ML. They have to build their strategy and intuition muscles around AI's potential and limitations, and constantly communicate these concepts to their community of stakeholders, teams, and end users.

There's a reason why AI PMs exist. The skill set you need to build machine learning models and craft solutions to AI-specific problems is different from the skill set you need to manage a product from vision to deployment. An AI PM's role is to translate the complex world of AI into a product roadmap; it's to balance technical possibilities with business goals and user needs. At its core, the work of an AI PM involves a lot of communicating. Whether it's explaining AI capabilities and benefits to non-technical stakeholders, aligning engineers with user-centered outcomes, or ensuring AI products are built ethically and responsibly, your work touches virtually every facet of your company. People have short attention spans. They have their own work to focus on! That means you're going to consistently reinforce ideas and concepts, continuously communicate the vision you've defined, and regularly give updates to all stakeholders. Communicating doesn't just mean talking in meetings with stakeholders, though there's a lot of that too.

You communicate through many channels: product documentation, blog posts, thought pieces, white papers, user stories, and changelogs. Your calendar is a way to communicate your priorities. You communicate regularly with end users and customers, directly or indirectly. You communicate regularly with your leadership team to make sure any threats or risks to the AI product strategy and vision are discussed. Your product marketing efforts communicate the value and viability of your product to your potential customers. The roadmap you settle on communicates which features and capabilities are making it into the product and which aren't. Those decisions communicate the strategy and vision of your AI product.

Communication is key! As an AI PM, you'll need to make sure you're aware of how your thoughts, ideas, words, and actions are serving your product communications. It's also a way to build trust with all the relevant stakeholders you'll regularly be communicating with.

The following are a few strategies for how to communicate effectively with all groups as you're managing your AI product:

- **Cadence**: Keep a recurring meeting with your key stakeholders. AI PMs can create a tiered list of stakeholders and decide how often you should be meeting with them. Don't leave strategy meetings with leadership, marketing, sales, and customer success ad hoc. While you'll often be communicating with devs, particularly for an AI PM role, you still need to keep this cadence with your other stakeholders strong. You *will* have things to discuss and the landscape of AI is often changing. This doesn't just go for internal stakeholders but for customers as well. You shouldn't be going long periods without hearing from customers. Scheduling recurring feedback sessions with customers and regularly sharing those data and insights with stakeholders is a big part of this.

  Here are some accompanying questions:

  - Are you having regular touchpoints with every department that you need to be collaborating with?
  - Are certain meetings too frequent or too infrequent?
  - Does your leadership team participate in product strategy meetings and are they engaged?
  - Do you have recurring feedback sessions with a diverse group of customers?

- **Transparency**: Don't keep a product team in an ivory tower. AI PMs need to be accessible because you're the mouthpiece of the AI product and often the one communicating priorities and decisions. Your roadmap shouldn't be a secret document that only your CEO and you see. It should be a reflection of your product strategy and vision. That product strategy and vision should be a reflection of your company's vision and goals. Everyone should be aware of how to understand your roadmap and where to find it, along with other pieces of product collateral and documentation. Being transparent about how decisions are being made and how things are being prioritized is also a way to communicate what the product organization values at any given time.

  Here are some accompanying questions:

  - Is your AI product roadmap publicly available and accessible to your stakeholders?
  - Does your product roadmap reflect your product strategy and vision?
  - Does it have buy-in from all major stakeholders you're collaborating with?
  - Does the roadmap outline how AI technologies, features, and capabilities will be integrated into your product over time?
  - Are there specific timelines for feature releases and updates for AI capabilities?

- **Updates:** Changelogs are a way for you to update devs on the customer side with the latest updates to your product. Blog posts and white papers are a way for you to update your greater market and potential customers on what makes your product special and effective. Newsletters, in-app alerts, and email lists are a way for you to regularly update your existing customers on product updates. Meetings with stakeholders and internal company-wide meetings as they relate to products are a way to update folks internally in your organization about what is happening with your product.

  Here are some accompanying questions:

  - Do you have a comprehensive strategy in place to handle all AI product updates internally?
  - Are your customers and prospects finding product updates through your communication channels?
  - Are you regularly checking marketing collateral, product one-pagers, and white papers to make sure outdated information is being updated to reflect the latest developments, particularly for AI features and capabilities?
  - How frequently is the roadmap reviewed and updated to adapt to market changes or user needs?

Communicating is such a core facet of PM work that it deserves its own category. Crafting a vision is one thing, but that vision is nothing without effective communication around it. It's not easy to articulate a clear vision and include everyone's two cents in it as well. Every additional perspective has the capacity to muddy the waters and dilute your message. These are all highly discerning activities.

Though the preceding section outlines strategies for effectively communicating your product vision and roadmap, you may have to be cautious about how often or how openly you communicate certain ideas over others. In some cases, it may be impractical to be an open book. Only you will be able to decide what degree of transparency, collaboration, and inclusivity you'll be able to handle. But more often than not, it pays to keep lines of communication open and to communicate your product vision, strategy, and roadmap in verbal and written form often. Probably more often than you think! In the next section, we'll discuss why this frequency and transparency matters more than most AI PMs let on.

The heart – Managing people and values

In the last section, we covered the more cerebral side of the AI PM role that deals with product vision, strategy, and communication. In this section, we will cover matters of the heart. The work of managing people is hard and AI PMs will have their fair share of personalities to work with. Between the stakeholders, teammates, customers, and prospects, that's a lot of people's hopes, dreams, wishes, and complaints to deal with. While AI PMs often don't have any direct reports or concrete authority over others, unless they're PM managers, directors, and so forth, they need to spread their influence to be effective at the work of product management nonetheless. The AI PM individual contributor role will need to be skilled at working with various functions like data scientists, engineers, and UX designers. Balancing their feedback and inputs into the product while making them feel safe and empowered to share their perspectives is crucial. Encouraging open dialogue and an open exchange of ideas limits bias and allows teams to build and work together in a way that fosters trust and collaboration.

Safety

Fostering a culture of psychological safety is important for any collaborative work environment. We spend more time with our co-workers than we do our own families much of the time, so working with people that we trust and feel safe with is imperative for all workers. The AI PM role is one that gathers inputs from all sides and finds a way to digest them all in a way that best serves their product. But what will the quality of that feedback be? If our stakeholders, collaborators, and teammates don't feel comfortable sharing their authentic thoughts, asking questions, expressing controversial opinions, or raising concerns without worrying about how others will react, the feedback won't be of much use. It will also alienate the people you need to build trust and work well with. An effective AI PM knows this well and helps create a culture based on respect, inclusivity, collaboration, communication, and trust.

This is easier said than done. A kind of hubris can develop in PMs, particularly now that AI is reaching a larger-than-life influence. This hubris can develop for many reasons. Because the role is so encompassing, the cumulative effect of having to balance so many perspectives, meetings, and resources can lead to a mental environment where you feel overwhelmed and underappreciated. Some PMs may feel resentment or a sense of superiority over others, leading to a downward spiral of trust and safety.

Remember that as an AI PM, you will likely absorb any criticisms, disagreements, and resentments stakeholders voice when you're collecting feedback from them. Things won't always be perfect. Sometimes we're disappointed when a feature we really wanted doesn't make it into the product, or when a decision is made that isn't in our favor, whether we're an AI PM or a key stakeholder. Whether it's a salesperson who really needs a piece of functionality for their biggest customer to buy or a machine learning engineer who disagrees with a modeling decision one of their peers made, they will often go to the AI PM first. Over time, the cognitive load from these factors can add up.

Our top priority should always be ourselves because we can't create a culture of safety and empowerment without first feeling it ourselves. I urge all AI PMs to establish strong self-care practices. You will find yourself empathizing with people time and time again. Absorbing too much of this will leave you feeling depleted. Absorbing too little will make it hard for you to establish authentic relationships with the people you will rely on to build and ship your product. You're at the nexus of a big ecosystem and your own health will radiate into that ecosystem.

Humility can't exist when we feel that certain people are more worthy or capable than others. It comes from the understanding that everyone has a part to play in making something great together based on their skills and inputs. As an AI PM, you'll have to balance a lot of different skills and knowledge to perform your role effectively. But it will largely be other people who will carry out the "work" itself, for the most part. Your counterparts in machine learning and development will be doing the development work to build and ship your product. Your coworkers in marketing and sales will be communicating about your product based on your collaboration. Your leadership will be evangelizing your product when they go to conferences. No one does anything alone. You need them and they need you!

The following are a few strategies for how to effectively build trust and safety with all groups as you're managing your AI product:

- **Human-centric approach**: A lot of people these days are concerned about what AI adoption means for their roles and job security. AI is also contributing to a lot of content being generated and personalized for them as well. What this means is that AI can contribute to a culture that actually negates the human elements. When you're working closely with people or starting a new working relationship with them, particularly as an AI PM, treat people like humans first. What that looks like is getting to know them; showing an interest in who they are outside of their role. What are their hobbies? Do they have a family? A side hustle? Showing an interest in who people really are and making an effort to build an authentic relationship means a lot to people. We're more than who we are at work and showing that people can be safe with you upfront will only help you maintain trust down the line.

  Here are some accompanying questions:

  - After you assess your relationships with your closest collaborators, can you say you know the person behind the role?
  - Are you also letting yourself be seen and appreciated for the person you are?

- **Inclusivity**: Being an AI PM isn't about coming up with all the best ideas on your own, it's about creating a safe space for all perspectives to come together and be heard. None of your key stakeholders should be left out of the product strategy and vision. Your roadmap should not be a surprise to anyone. You have to build trust and buy-in with all stakeholders if your roadmap is to be accepted and that begins with including them early and often. When it comes to AI products, inclusivity is critical because limiting perspectives and viewpoints means bias may be inherently baked into how your AI product operates and makes determinations. Varying viewpoints can help identify potential biases, ethical considerations, and unforeseen impacts on your AI product and its users. No decisions should be made in a vacuum; that's not what leadership is about. Collaboration is a cross-functional sport.

  Here are some accompanying questions:

  - Do you feel you're accepting perspectives from all departments that impact your product?
  - Can you confidently say your peers feel heard?

- **Objectivity:** The world of an AI PM can be hectic and fast-paced. It can be nice to have an objective way of vetting which features and tasks make it into a backlog or sprint and which don't. An effective way of managing the emotions of various stakeholders is establishing a way of assessing certain requests against their impact on customers, costs, urgency, or resource-heaviness. This is particularly relevant for AI projects, where the complexity and potential implications of features can vary widely. Some tasks are less feasible than others. Some are more costly or might be more time-intensive than they may be worth. If you're able to establish criteria for which features make it, and if you can be transparent about your product roadmap as we mentioned in the previous section, you can manage the expectations of others a bit better.

Here are some accompanying questions:

  - Do you employ any scoring methods when assessing or prioritizing AI features?
  - Have you communicated to stakeholders what methods are used for prioritizing your product backlog, roadmap, or epics?
  - Are you considering the ethical implications and potential biases of features when assessing their priority?

- **Modeling:** The best way to get buy-in into a culture you're helping create is by demonstrating it yourself first. When you're in meetings with stakeholders, you set the standard for how you want those meetings to go. If you're conducting a recurring meeting, workshop, or brainstorming session, find ways to embody the spirit of experimentation, curiosity, and commitment to your product. Even something as simple as admitting when you make a mistake or an assumption can go a long way. It shows you can still be capable and respected even if you're sometimes wrong. In order to be an effective AI PM, you have to be a credible source of truth for your organization. You can't foster credibility if you're not practicing what you preach.

  Here are some accompanying questions:

  - Do you feel you're in alignment with the cultural goals you've set for your product team?
  - Do you embody the values and strengths you ask for in others?
  - When was the last time you publicly admitted a mistake without shame?

- **Recognition:** Being an AI PM allows you a wide perspective into the contributors that most impact your product. As you allow more perspectives to influence your product, you'll naturally start to see people showing dedication, consistency, and achievement. This is an opportunity to use your vantage point to recognize those contributions. It makes people feel appreciated and it helps others buy into the culture you're trying to create. Just simply saying you promote a collaborative work culture isn't the same thing as embodying that collaborative spirit and recognizing it in others. It also helps reinforce the idea that you say what you mean and you mean what you say, which is usually in short supply. Be timely about it. Recognize and celebrate your peers when their contributions are making an impact on your AI initiatives.

  Here are some accompanying questions:

  - Do you remember the last time you recognized someone else's contribution to an AI project or feature publicly?
  - Do you remember the last time you were recognized for your work and how did it impact your motivation?
  - Can you incorporate recognition into your regular team meetings or communications to reinforce a culture of appreciation and gratitude?

Safety is integral for any relationship to succeed, let alone a company full of them. Given the scrutiny on AI products today, we can't afford to have companies investing in AI product teams without seriously ensuring those teams are building in a way that feels safe to internal stakeholders, customers, and prospects alike. If you're not open to new ideas or if you can't effectively prove the values you're leading with, your PM style won't be authentic to your stakeholders. This will undermine any efforts, campaigns, plans, and roadmaps you set out to implement.

Empowerment

Companies of all kinds will need to think critically about their AI strategies if they want to remain competitive in the AI landscape we find ourselves in. A big part of creating a culture around AI is getting clear on what kind of AI company they want to be. Rethinking how your AI product will transform your customer experience and how you can communicate that will help your product marketing efforts and it will also orient your internal teams toward a common goal. Empowerment isn't just about making sure the data scientists, machine learning engineers, product analysts, or UX researchers are comfortable with various facets of AI. It's about equipping the other stakeholders in your orbit to think critically and constructively about how AI will be incorporated as you're building. This will allow you to be able to fully appreciate the feedback and collaboration you get from your peers when you can help them to more fully integrate AI concepts.

The following are a few strategies for how to effectively empower your teams and stakeholders to embrace AI as you're managing your AI product:

- **Agility**: Maintaining an agile mindset is important for the AI native product. You won't necessarily settle on models and hyperparameters that work best for your organization, data, and customer expectations without having an experimental, nimble approach. Staying flexible and allowing the process to iterate as it needs to is important. Often, it's unrealistic presumptions and overblown expectations that lead to disappointing results with AI, so remaining agile allows for product and development teams to build without letting expectations weigh them down.

  Here are some accompanying questions:

  - Are your recurring feedback sessions structured to encourage fresh ideas and prompt new product iterations or AI features?
  - What was the most impactful product iteration that significantly improved user experience that came from insights gained during UX research?
  - How do you integrate lessons learned from experiments and user feedback into your AI product development cycle?

- **Ethics**: AI products are powerful and this is precisely what makes them a potential risk factor for the organizations that develop them. All major stakeholder teams should feel empowered to bring up concepts regarding ethics, data privacy, legality with local laws, and mitigating algorithmic bias when it comes to building AI products. These considerations shouldn't only be the domain of the product and ML teams but they should be considerations people from all teams should feel empowered to voice. AI PMs can play a foundational role in this by supporting training efforts to get their organizations as AI and data-fluent as possible.

Here are some accompanying questions:

- Do your stakeholders feel comfortable voicing ethical or safety concerns during your recurring meetings?
- How does your organization facilitate open discussions about ethics, data privacy, and algorithmic bias in AI product development?
- Does the company encourage employees to upskill or receive AI training and other efforts to promote data and AI literacy?

- **Innovation**: Once you get into a day-to-day routine, the inertia of a repeatable sequence sets in. Breaking out of this can be hard for organizations that have tight deadlines or inundate themselves with milestones. It can be tempting to do this. Nothing makes someone feel like they're useful and capable quite like a checklist. But often, we fall into the trap of creating urgency and scarcity in our efforts to feel useful. Try to create space in your organization for experimentation to thrive. Perhaps you can bring a small group of people together or encourage them to spend 10% of their work hours exploring new ideas. Or you could set up an AI innovation lab and let that be a day once a month where your research, AI, data, and product teams can collaborate on innovative new product offerings. Consider fostering a culture of innovation by establishing dedicated time for teams to brainstorm and develop new ideas. We're still early in AI adoption and organizations that prioritize this stand to gain a lot. Doing so will help promote a culture of curiosity, which will empower your organization to dream up new possibilities that otherwise wouldn't exist.

  Here are some accompanying questions:

  - When was the last time you and your team felt you had the brainspace and freedom for a truly out-of-the-box brainstorming session focused on AI innovations?
  - Does your leadership team see innovation as a worthy investment in the future of the company?
  - How are new ideas and experiments prioritized within your AI product development cycle?
  - What structures are in place to ensure innovative ideas generated during brainstorming sessions are captured and effectively acted on?

- **Boundaries**: In the world of AI, there's always something new and exciting to read or discover. You could be working around the clock. But promoting a positive work culture has to be a priority for PMs. You shouldn't be sending communications out after hours or working during your vacations. We lead by example; doing so communicates to your peers and stakeholders that this behavior is acceptable and that they should do it too. Burnout is hard to recover from and as an AI PM, you should be avoiding that at all costs. We aren't machines; we're biological creatures that need rest and strong boundaries if we're going to take on challenging roles. You won't be effective for anyone if you can't maintain boundaries for yourself and nurture your well-being. Doing so sends a powerful message to others that they should too. The organization will benefit tenfold from having an energized, engaged workforce.

Here are some accompanying questions:

- Do you feel overwhelmed by your responsibilities or are they a source of excitement and motivation?
- Does your work ethic make others feel supported and does it encourage them to maintain their own boundaries?
- What strategies do you use to disconnect and recharge?

Working in an organization that prioritizes and invests in its employees is empowering. We should all be looking to work for organizations that value and model this behavior because those are the organizations that benefit from an engaged, dedicated workforce. When people feel empowered, they feel big and bold. It gives them the energy to push forward with love and without fear.

The guts – Managing the rest

In the world of an AI PM, managing "the rest" encompasses the resilience and adaptability to navigate the unpredictable landscape of technology and market demands. Responding to challenges as they arise, proactively identifying risks and opportunities that impact your product's success, and managing the day-to-day functional operations of an AI product is a heavy load. AI PMs frequently need to use data to inform decisions, iterate quickly based on user feedback, and stay on top of industry trends and competitor movements to remain effective. Cultivating a mindset of experimentation and embracing mistakes is a stepping stone to innovation and sustainability when it comes to managing AI products. As with the machines that learn from past mistakes, so do we. Building a gut instinct for resilience and adaptability will help manage the complexity of AI products in a way that keeps the team aligned and motivated for the long haul. Let's go over various areas we will manage over the course of our AI PM work that will help us approach challenges with confidence and creativity:

- **Data**: All AI products need to have an effective data management strategy to function properly as AI products. Understanding elements that relate to data collection, cleaning, preprocessing, scraping, and privacy will ensure that you're effectively managing the data side of your AI product. Data management issues can eventually have an impact on data quality and significant issues there will render your AI systems useless. Improper management and maintenance of data could open you up to legal risk as well, particularly if you're not adhering to data privacy regulations. Over time, we're likely to see more data privacy regulations and legislation so make sure you're prioritizing this work.

  Here are some accompanying questions:

  - How do you ensure data quality, integrity, and governance throughout the life cycle of your AI product?
  - Have you examined all relevant data privacy, collection, and processing regulations that are applicable to your AI product?
  - Would you be proud to talk publicly about your data management practices?
  - What strategies do you employ to monitor and improve your data management practices?

    - Have you established a framework for ethical data usage that includes considerations for bias and fairness in your AI models?
    - What process is there in place to incorporate feedback from data scientists and ML engineers regarding data management challenges/opportunities?

- **Centralize**: A lot of communication happens asynchronously when managing AI systems due to the cross-functional nature of AI development, so it's great to have a central repository where people can find ways to collaborate and communicate that way. Tools like Confluence, Notion, Jira, Trello, and Asana all serve as a knowledge base or a place to come together and ask questions relevant to certain areas like tracking AI/ML model development, managing data pipeline progress, and user research insights. They also communicate what product areas are being worked on at which time. The most challenging part of this is getting people to use and rely on them consistently but once they do, you're reducing the need for synchronous face-to-face updates. This also means that when you do meet synchronously, you're making the most of the time you do have on more substantial, in-depth, and strategic conversations.

  Here are some accompanying questions:

    - Are the AI/ML and data science teams effectively documenting model experiments, metrics, and decisions within your centralized tools?
    - How do you make sure communication breakdowns between AI, data, and product teams are minimized, especially when experimenting with new models or datasets?
    - Are there tools you've invested in that your teams don't like or don't want to use?
    - Can the tools you're using help automate routine communication like reminders or task updates to reduce manual efforts for your teams?

- **Training**: Often people don't end up using certain tools or adhering to certain practices because they haven't been properly trained for them. Not everyone is going to be excited about adopting a new tool or learning a new way to do something. Inertia is real and most people want to streamline their work as much as possible. Learning something new can often be an impediment to that, so encourage your team and the stakeholders around you to join training sessions to better understand PM processes and tools. Training is also a good way to keep various teams up to date on new trends, AI best practices, or even new AI tools that have recently come on the market.

  Here are some accompanying questions:

    - Is there a way for you to gauge people's interest or satisfaction with certain tools, particularly those that have AI capabilities?
    - Are any tools in the current tech stack too difficult, time-consuming, or cumbersome to use effectively?
    - Do non-technical stakeholders receive training on AI concepts relating to data science, ethics, bias detections, fairness, and model interpretation to bridge any gaps?

- **Feedback**: Scheduling time with various teams like data science, ML engineering, and UX and getting feedback from them regularly is a great way for you to make sure stakeholders on other teams feel their perspectives are wanted and appreciated. AI products do best with continuous iteration so regular insights from stakeholders will refine your AI native product in meaningful ways. In the previous section, we mentioned cadence with stakeholders. This is applicable here as well. You can make sure that receiving feedback is on the meeting agenda for those regular meetings you establish with key stakeholders. If you don't keep a regular cadence with certain teams, you can create reminders on a monthly or quarterly basis to book time with them so that you're including all necessary perspectives. Incorporating regular feedback sessions into your process can mean the difference between a stagnant AI system and one that evolves by meeting the needs of your customers, users, and market conditions.

  Here are some accompanying questions:

  - Is there a plan in place to translate technical feedback you receive from your data scientists and ML engineers into actionable steps for your AI models, features, or capabilities?
  - Are you receiving feedback from a variety of sources like sales or marketing to ensure your AI-native product remains customer-centric?
  - How often are you engaging regularly with external partners or AI vendors for feedback on the AI tools or platforms you're using if they're third-party?
  - How often is feedback regarding model accuracy, bias, or performance fed into your product roadmap, epics, or backlog prioritization?
  - How are you balancing technical feedback related to AI/ML model tuning with the broader concerns on customer impact when you're making edits or updates to your product roadmap?

- **Workshops**: This one seems nice to have but it's really not. Creativity thrives in play, so organizing occasional workshops to explore ideas, new strategies, and perspectives is a great way to identify new product areas, features, and gaps in your current market, particularly when managing an AI product, where use cases are still evolving. AI PMs can hold core product workshops that focus on model selection, data collection, or explainability, as well as marketing workshops, sales workshops – and even product strategy and market analysis workshops to try and broaden your AI product's use cases and offerings. Make it fun, low-stakes, and collaborative, and see what magic comes out of them. As long as people feel safe, free, and appreciated, something will. It will also strengthen trust between teams and reinforce the benefit of a viable AI product team.

  Here are some accompanying questions:

  - What were the most innovative AI-driven features, solutions, or ideas that have come from your workshop sessions?
  - Have you used workshops to identify new AI use cases, datasets or model approaches that hadn't been considered before?

    - Is there a strategy in place to encourage non-technical team participation in AI workshops to make sure broader perspectives are being welcomed into your AI product development process?
    - What was most surprising for you from the product workshops you've conducted?

- **Testing:** AI products should be thoroughly tested and vetted. We've covered this topic extensively in other chapters but it's an important part of the day-to-day work of managing an AI product. Unit tests, integration tests, user acceptance tests, and bias tests have to be a recurring part of the AI product development cycle. The complexity of AI systems makes testing more crucial because model performance can degrade over time and changes in new data could introduce unexpected issues. The more products are tested and proven, the more you'll build trust with your customers which will, in turn, help you to convert prospects as well. Word spreads and the more you can prove you're rigorously testing your AI products, the more it'll benefit your product performance and brand identity.

  Here are some accompanying questions:

    - Have any issues or product anomalies in the AI system not been addressed?
    - Have your AI models been tested for edge cases, bias, and fairness?
    - How frequently are models being evaluated for model, performance, or data drift and degradation?
    - Are there areas of your AI product that still need to be tested or monitored for performance?
    - Is there a continuous testing strategy for retraining models or updating data pipelines?
    - How are test cases handled when models perform inconsistently or demonstrate bias in specific demographic groups or regions?

We've covered a lot of ground in this chapter. We discussed overarching elements of the AI PM role, which included maintaining alignment across your AI product development, as well as managing the values and motivation of the people that will be a part of your product organization. We also discussed the more routine elements of managing the day-to-day AI product operations. We'll pull these themes together with the help of our case study to see them contextualized into an applied example.

Case study

Waterbear started as an AI-native company and a big part of their AI and product strategy was centered on how to empower teams to embrace the promise of AI as it related to mental health support. The tool was never intended to replace a human therapist in any way; it was created as an accompanying tool for anyone who wanted to know themselves more deeply. The purpose of the product was to help users understand their own psyche better. Even when we write in a journal, we miss out on insights and trends that emerge in our writing. Using NLP and the power of LLMs was a practical way of helping users understand their own dreams and fears better, to see them for what they are, and to use that knowledge to unlock their end users' biggest goals.

This was a foundational part of how the company started; the founders knew they wanted to use AI in a psychological and mental health context. The company mission was simple: "*Know thyself with the help of AI.*" It was often a part of their marketing campaigns and pitch deck. They wanted something that resonated with a deeply unsettled population, especially women who faced challenges in balancing work and personal life. This was particularly relevant in a society where women's rights were under scrutiny and they were increasingly taxed without full bodily autonomy. For this reason, Waterbear crafted the following product vision statement to center its AI products around: *"Waterbear uses AI to empower individuals globally, offering insights that unlock deeper self-understanding and personal growth."*

Waterbear wanted the focus of their product vision statement to emphasize empowering individuals to convey a sense of agency and personal strength, as well as the global nature of the problem they're looking to solve. Highlighting insights and deeper self-exploration reflects the core purpose of the product, which is to fuel significant personal revelations in a way that's self-directed. The aspirational part of the statement resonates from the focus on personal growth, which is a journey that never really ends and can look different for everyone. It's also concise, clear, and easy to remember and communicate across channels.

The ramifications of that, exacerbated by a mental health crisis from the pandemic, created an environment where Waterbear's founders were ready to go all in on their product vision and start to build out the team for Akeira, their flagship product. The founding product team for Akeira wanted to focus on the core competencies of the AI PM role: AI literacy, communication, market research, user research, strategic thinking, analytics, leadership, customer engagement, and collaboration. In that spirit, one of the first things they did was conduct user research trials and begin the work of scouting therapists that they could start working with. Once they had this team, they began the iterative work of trial and error to discover what combination of factors would result in the best product experience for their users.

The product (Akeira) aligned with the company's vision and mission. It helped women specifically know themselves and empowered them to reach their goals faster. Waterbear knew this was only the beginning, so the goal was to capture more groups of users and offer them an AI of their own to connect with eventually. They knew they wanted to learn from the experience with Akeira, and to get her to a point where they could launch an MVP that they would be proud to share with their beta users. Once they did, they knew their product team and greater organization would put in the dedication to serve a loyal customer base. Post-launch, a beautiful journey of discovery began.

The following is a list of the most significant actions from the founding product team as they began this journey, along with their impact:

The head – managing alignment		
	Action	**Outcome**
Vision	Officially launched their product roadmap, along with the MVP, after they received approvals from all major stakeholders at the company.	This helped ensure they had alignment and could confidently move forward with their plans.
	Set up the metrics, KPIs, and OKRs to best support the growth and trust of Akeira.	This ensured all measures of success were in service to their goals for Akeira and Waterbear.
Cadence	Set up recurring meetings with sales, marketing, support, leadership, ML, and data teams.	This made it possible for ongoing efforts to be reported on and discussed with all impacted teams.
	Set up recurring feedback sessions with their beta testers.	This ensured they were keeping their customer voice central to their research efforts.
Transparency	Set up a monthly customer-facing changelog section of the Akeira app where updates to data practices and security had their own section.	This allowed the customers who were more technically inclined and vocal in their virtual community to connect more deeply with the product and ML team for Akeira and to share those insights with the broader community and serve as brand ambassadors. It also helped reinforce a brand identity around trust in data and AI.
Updates	Made company-wide announcements anytime there were any changes to their roadmap.	This absorbed the shock of any important updates that might have fallen through the cracks.
	Set up a recurring monthly "office hours" meeting to give all departments updates on what was built and shipped for Akeira for the month.	This served as a catch-all for any questions anyone at the company might have from a product perspective.
The heart – managing people and values		
Human-centric approach	Conducted monthly "get to know your X" sessions where various team members were randomly matched to learn about each other personally.	This allowed team members to get to know each other on a personal level irrespective of how often or deeply they work together professionally. This approach supported the idea that although you might not work with someone very closely today, you might in the future.

Inclusivity	Implemented a bias review committee that met quarterly to address the inclusivity and ethical considerations of Akeira.	This allowed a structural approach to maintaining trust in Akeira that was well balanced and represented at the company. Stakeholders from various departments, as well as trusted super users, were invited to participate in discussions that allowed for diverse viewpoints to be considered.
Objectivity	Employed a public-facing feature evaluation matrix.	This allowed features to be evaluated based on customer impact, cost, urgency, and resource requirements, along with caveats when concessions or overarching decisions that went against objectivity were made in efforts to build transparency and trust across teams.
Modeling	Created a Slack channel called “hands up” that anyone could rely on whenever they started to get overwhelmed by their workload and the leadership team rewarded those who posted on the channel.	This showed everyone at the company that there’s no shame in feeling overwhelmed. Often, it was members of the leadership team that needed it and used it the most and others followed suit once they realized it was safe to do so.
Recognition	Created a Slack channel called “shoutouts” to publicly recognize and promote the contributions of various members of their organization.	This kept the momentum strong post-launch and made people feel seen and valued as they poured their energy into the strength of Akeira.
Agility	Set up the model experimentation lab where sandbox environments were set up to test different model versions for Akeira.	This encouraged the team to test different NLP models and hyperparameter settings in a safe environment for conducting A/B tests and comparing performance metrics and iterate quickly when performance improvements did jump.
Ethics	Recurring bias audits were conducted across all major user personas to gauge potential drift.	Rather than monitoring performance across the platform as a whole, bias audits were conducted around each major user persona for Akeira, ensuring bias was reduced across user types and not across all users as a whole.

Innovation	Insights and major breakthroughs from their AI innovation lab were consolidated into monthly Innovation Days.	Though the data science and ML engineering teams devoted a percentage of their working hours to the innovation lab on a weekly basis, Innovation Days were a way for these teams to communicate major findings with the rest of the company and functioned as a built-in AI training for non-technical stakeholders to get better acquainted with AI and NLP.
Boundaries	A No-After-Hours-Communication policy was implemented to reinforce healthy work-life balance.	All employees were instructed to refrain from sending emails and messages, or scheduling meetings, after 6 PM local time to maintain clear boundaries between working hours and personal time off.
The guts – managing the rest		
Data	Data pre-processing protocols customized to the needs of Akeira were put in place to tokenize and normalize user inputs, which were often long-form.	This allowed Akeira to tokenize inputs (break them down into smaller units like words and phrases) so that words, meaningful components, context, and sentiment could more effectively be ingested by Akeira's AI system.
Centralize	Built a product knowledge base where every feature of Akeira's has a dedicated Notion page where all product requirements, information, documentation, style guides, and feedback can be centralized.	This allows for all Akeira documentation to be searchable and accessible to all internal teams. It also allows for all questions, concerns, and recognition to be centralized.
Training	They conduct a quarterly training session on the efficacy of LLMs and potential ethical considerations that come up when building with conversational AIs.	This proves the company's dedication to building trust with its dedicated user base because it invests in its employees better understanding the tech that powers Akeira. Even if they aren't technical, Waterbear feels their product experience is best when everyone at the company understands how it works.
Feedback	Monthly feedback debriefs focused on user feedback from the Akeira community, social media, and in-product surveys were conducted.	This allowed the opportunity for data science, ML engineering, and UX teams to regularly review feedback coming from the wider Akeira community and for the AI PM to incorporate that feedback into the active user stories for features that are being worked on and planned.

Workshops	A "workshops" Slack channel was created to gauge interest in various AI product categories and a Google Calendar was regularly updated with the monthly workshop.	Building a culture of continuous learning at Waterbear meant that a regular calendar of upcoming workshops was circulated company-wide. Staff favorites were also highlighted as recurring workshops to meet the demand where applicable.
Testing	Akeira super users were encouraged and incentivized to become a part of the Super Testers Squad, a group of external users dedicated to maintaining integrity in the platform.	Keeping super users and super testers squad communities separate allowed for feedback to come from two separate groups and separated users based on their level of dedication to maintaining Akeira's quality standards. Users were also recognized and rewarded based on their level of dedication. The Super Testers Squad was in regular communication with internal testing teams and bolstered trust in the Akeira platform across all users.

Table 11.1: Significant actions from the product team

This is by no means an exhaustive list, but it is a list of the practices that contributed most to the success of Akeira's founding AI PM team. Most of the items on the list were actions that contributed most to the emotional and cultural health of the product organization. That's because the founding product team knew that when they supported their stakeholders, customers, and prospects as best they could, they'd find success with their product. They already knew they had a great technical team. There was already a lot of trust built with the ML and data teams that supported the product organization. This meant that many of their efforts and actions went toward ensuring the rest of the world around them knew it as well.

Given the success and adoption of Akeira, it's clear the founding product team's intuition and values were in the right place. The reason why they focused on their values and building trust was because Waterbear's founders already knew there would be pushback with a conversational AI. The market is saturated with conversational AIs and there is already a lot of distrust when it comes to AI companies' data storage practices. Even if consumers weren't inherently suspicious or aware of these things, it still felt like a long shot that the average consumer would bare their soul to an AI. This is why so many of their efforts early on were focused on over-communicating and building a brand identity around trust and security.

Summary

In this chapter, we discussed how managing alignment across leadership, stakeholder teams, and day-to-day operations is crucial to the success of an AI PM. We learned that establishing a clear product vision, strategy, and roadmap that reflects the company's overall goal is necessary because the use of AI technology must address user-centric problems through cross-functional team collaboration. Misalignment can lead to inefficiencies that are too costly in such sensitive environments. AI PMs must also navigate the volatile nature of technology, market demands, diverse personalities, and competing priorities to do their work effectively, as we learned in this chapter. We further learned that by cultivating a curious mindset and an open culture of collaboration and experimentation, AI PMs can hedge their bets against the many sources of entropy that can influence their AI product development cycle.

This chapter concludes *Part 2* of this book on the AI native product. So far, we've had a chance to understand the AI-native product from a team, tech stack, productization, customization, and positioning perspective. We've discussed the many facets of what it takes to build products natively with AI and took a deeper dive into benchmarking, growth hacking, and storytelling for your AI native product. We've also explored the most important product design and product management considerations that are tailored to the unique challenges that face AI PMs.

In *Part 3*, we'll explore what it looks like to build AI capabilities and features into existing traditional software products and meet a new case study example that will demonstrate the concepts we discuss in *Part 3*, much like we did in *Part 2*. We'll start *Part 3* with *Chapter 12*, which will focus on the rising tide of AI and how product teams can prepare themselves for the demands of working with AI when they're already established.

References

- *Pinterest vision statement*: `https://www.pinterestcareers.com/our-life/`

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

Part 3

Integrating AI into Existing Traditional Software Products

Many companies will embark on the journey of integrating AI and ML into their existing traditional products because of the potential competitive advantage and strategic influence AI can provide across industries.

The third part of this book will focus on evolving existing software products that don't currently use ML or DL to leverage AI. In the previous part, we discussed building an AI-native product and how the process unfolds in a way that's optimal from a data, company, and strategy perspective. Using this same lens, we will now compare and contrast how this process unfolds for products that don't currently leverage AI. By the end of this part, we'll understand the impact AI is having on individual companies and the industries they're a part of. We'll also see how product design and product management are impacted with the introduction of AI features and capabilities for products that have already been launched prior as traditional software products. Similar to *Part 2*, we will contextualize these concepts through different case studies in this part.

This part comprises the following chapters:

12

The Rising Tide of AI

With this chapter, we begin the third part of this book, which will focus on evolving existing products that don't currently use **machine learning** (**ML**) to leverage **artificial intelligence** (**AI**). We will refer to these products as non-AI-nnative products. In the previous part, we explored how to go about building an AI native product and how the process unfolds in a way that's optimal from a data, company, and strategy perspective. Using this same lens, we will now compare and contrast how this process unfolds for products that don't currently leverage AI.

More specifically, this chapter will re-introduce us to the concept of the Fourth Industrial Revolution. It will serve as a reminder and a blueprint for those of us working as entrepreneurs or PMs for products and organizations that are ready to fully embrace AI. Because this chapter introduces *Part 3* of this book, which focuses on integrating AI into traditional software products, it serves as a good reminder of what companies stand to gain. We will be reinforcing the importance of the coming shift toward AI and what it means across industries. For those who are already at a place where they can embrace AI, I hope to offer a plan you can be confident in as you move forward toward evangelizing AI in their products and wider organizations. For those who are still skeptical, I hope it galvanizes you and offers a clear vision of what's to come so that you can start the work needed to prepare for such a shift. For those who are AI skeptical, I hope to offer a glimpse into what a future without AI adoption looks like for your company's success.

As we continue in this chapter, the first section will go into what the AI evolution will mean to businesses and products in order to best communicate the competitive advantage of embracing AI at the industry level. We will then go over accessible options for adopting AI for companies looking to capitalize on the AI adoption wave in the second section. The final section will focus on the mindset and attitude for AI adoption at the PM level, as well as how PMs can communicate AI adoption across their organizations. I'm honored to be a part of this journey with you, and I hope the coming words offer you solace in a landscape of misinformation and disinformation when it comes to AI and its adoption.

In this chapter, we will cover the following topics:

- Evolve or die – when change is the only constant
- Structural changes in the Fourth Industrial Revolution
- Fear is not the answer – there is more to gain than lose (or spend)
- Anticipating potential risks
- How LLMs are evolving

We will contextualize this with the help of a case study at the end of the chapter.

Evolve or die – when change is the only constant

I like referring to the evolution that AI will require of all industries as a *competitive imperative* because the language we use to describe this paradigm shift will need to be more urgent if it's going to adequately prepare people for what's to come. Each industry is a universe in and of itself, and once companies within industries start to see success with AI adoption, this adoption will accelerate. I've mentioned before that product management is an inherently commercial role because, as a PM or leader, you are tasked with the commercial success of your product. The competitive edge and marketing splash that can come from promoting an AI product are undeniable.

But beyond marketing, AI allows all industries to advance in more tangible ways. All economic and social areas will change with AI, from healthcare to tech, to education and government. The way industries relate to raw materials and plan, predict, and build products will change with this transition to AI adoption. The **World Economic Forum** (**WEF**) defines the industrial revolution we're currently in as *"characterized by a fusion of technologies that is blurring the lines between the physical, digital, and biological spheres."* I love this description because it places AI and ML firmly in the middle of these spheres. In order to account for the physical, we need a way to translate the physical world into data. This data could be about anything in the world beyond us, or it could be about ourselves. The collection and processing of this data will inevitably interact with AI and ML to some degree in order to be actionable in some way, and once this ball gets rolling, we humans, in our infinite creativity and inspiration, will come up with new combinations and expressions beyond the use cases we see today.

The way our lives are changing in seen and unseen ways affects many technological areas, not just AI. We're currently seeing massive breakthroughs in these areas, and the WEF goes on to say, *"the possibilities of billions of people connected by mobile devices, with unprecedented processing power, storage capacity, and access to knowledge, are unlimited. And these possibilities will be multiplied by emerging technology breakthroughs in fields such as artificial intelligence, robotics, the Internet of Things, autonomous vehicles, 3D printing, nanotechnology, biotechnology, materials science, energy storage, and quantum computing."* This is at the heart of what makes this such a privilege to witness in our time. There's effectively a technological renaissance happening, and we're slowly witnessing this renaissance unfold in real time every time we read the news, download a white paper, or find a study. I'm constantly reminding myself how exciting the current technological advancements are considering that 15 years ago, I didn't have a smartphone.

Even beyond the areas where we're experiencing breakthroughs, casualties or negative externalities, such as job losses, accompany all this progress. The automation of work will replace tasks and jobs that are routine, and almost all jobs have some routine to them. This is inevitable. It's also why some fear the age of AI because they're worried AI will replace their jobs. After all, if a company can get things done with automation, by buying a product, or by using some form of AI to replace a human worker, then why wouldn't it? There are people who work in operations, procurement, or finance whose entire job is to look for ways to save their companies money, and even those people could find their own jobs automated! Eventually, the shift will come, and in the meantime, we'll be reading about it one headline at a time.

But what's more likely is that AI will be used to assist just about all workers with their jobs in some capacity. I don't mean to minimize the collective loss that will arise with AI. There are jobs in data entry and manufacturing that will never come back, and you could argue they shouldn't. But AI could make all these debates a moot point by automating the very jobs we're looking to boost with a higher minimum wage. We've heard arguments that AI will replace jobs that are unenjoyable, mind-numbing, and oppressive. We've also heard arguments that AI will allow humans to capitalize on their curiosity and make even more money working jobs that place a premium on their creative, complex problem-solving, and critical-thinking skills. This assumes most people want that and don't want an easy job they can do day in and day out and live a peaceful life free of complex thoughts. Those who are in the know and can anticipate the way their jobs will change because of AI can better prepare for it, and they'll probably make more money because they're embracing this change. Others may find their jobs being automated out of existence. This is all but guaranteed.

Beyond all the doom and gloom, there's an upside. Many of the jobs people will be doing in the future are not yet known. According to a report by Dell and the **Institute for the Future** (**IFTF**), 85% of jobs that will exist by 2030 don't exist yet. That's incredible. What's perhaps most incredible is that this is all being discovered by all of us as time goes on. Just 15 years ago, we'd be confused if someone told us they were a social media manager. Even throughout my own career, I have seen what was previously called an account manager role shift and morph into a more sales-focused account executive role or a more performance-focused role such as customer success manager. The terms data analyst and data scientist also entered the zeitgeist not long ago. The function of both these roles has existed for some time, but the formalization of those roles and what they mean at a certain point in time is a new development. In the near future, we might be hearing about tele-surgeons, data brokers, drone managers, and VR technicians more often than we do now.

The evolution that AI adoption will require of us will have positive and negative effects, as with any technological adoption. Can we say we're all the better for it now that we have electricity running through the walls of our schools, offices, hospitals, and homes? How can we quantify the social and economic net benefit of such a massive change to how we live and work? How could the contributors of electricity have known that, eventually, their breakthroughs would power the internet and AI? If you could time travel back and tell them, they wouldn't even understand the gravity of what you were explaining to them in the first place. It's important to keep an eye on the risks.

We'd love to see governments and companies working together to embrace this technological shift and to see major sectors of industry recognize a need for skilled workers that can help them build the next generation of their products and nurture this wave of workers with paid internships, training programs, and leadership programs that will bring more workers into fulfilling careers with a future. I do a lot of speaking engagements addressing upskilling and career transitioning, and I repeatedly remind my audience that their interest, dedication, and commitment to upskilling are immensely helpful to society at large. Data analysts, data scientists, and ML engineers are highly sought after, and there's a huge skill gap. There are not enough skilled, trained, and experienced workers in these fields for the demand that's coming from companies. If you take on this learning journey, you will eventually find success, and you'll be rewarded with a lifelong career that will keep your interest. We can only hope that our larger institutions will see the upside in helping individuals with the burden of upskilling and invest in making that a smoother transition because it truly is a social and economic investment that's in everyone's best interest to be sponsored.

Changes in the Fourth Industrial Revolution

As a company and leadership begin the process of opening up to the capabilities and promise of AI, it's important to understand the full spectrum of opportunities AI can potentially offer. It's important for your company and leadership team to fully understand the scope and competitive advantage AI adoption will mean for your business. Product functionality is not the only area AI helps with.

Your strategy will also define how you approach AI program development. In the following sections, we will briefly explore the expected changes in terms of both culture and structure, followed by a discussion of the major areas of investment with regard to AI strategy, so you can choose the best course of action for your business. My recommended best practice would be to use these cumulatively. If you start with a consultant and escalate from there, you'll be setting your business up for success with regard to AI adoption and strategy.

Cultural and structural changes

Because AI is so data-hungry and progress is derived from performance benchmarking, you first and foremost have to have a healthy *data-driven culture* throughout the organization. You can't really see how AI will help you if you don't have a consistent baseline. Having metrics and KPIs in place to track your performance before you implement AI projects and AI features in your product is crucial for measuring and communicating success. Building this culture must come from the top down and I invite all leaders reading this to invest in building that culture to best see successes from AI adoption.

The second cultural influence I want to stress in this section is to have a *willing, open, and curious culture* around AI. According to Harvard Business Review, 80% of AI projects fail, a number that's almost double the rate of IT projects ten years ago. A big part of this is organizations not being appropriately set up for success regarding AI adoption. Inundating your data science, ML, and product teams with pressure and stress is not the way to set them up for success. The ability to iterate, experiment, and take some risks with regard to AI so that your teams feel safe to try various use cases is crucial to giving AI the best chance it has within your organization.

It's important to maintain a healthy skepticism and temper expectations with your AI activities until you experience early successes. It's even more important that you celebrate your early AI successes loudly and proudly across your organization when they do arise.

It's hard to overstate the gravity of what AI adoption will mean for all industries and all job roles, and the delineation between *technical* and *non-technical* roles will start to change as well. Right now, AI is mentioned in business articles for the most part as a rising trend or wave, but this wave is quickly turning into a tsunami. In order to stay competitive with their peers, all companies across all industries will find themselves scrambling toward the digital transformation of AI. As more and more companies do this and successfully advance toward accomplishing AI adoption, we will also be seeing more demand for *data-centric roles* simply because most products, internal operations, and client discussions will evolve along with the AI adoption strategies of companies.

We're also already starting to see **automated ML** (**autoML**) companies and offerings starting to grow. Companies creating autoML products such as DataRobot and H2O.ai are starting to pop up more and more. This will allow anyone, even those who don't actively work as data scientists and ML engineers, to be able to use, tune, test, and deploy ML models for performance. This means that over time, AI adoption will become more accessible to those who choose to take the red pill.

This also begs the question: *What will happen to data scientists and ML engineers*? Demand for these roles isn't projected to go down over time. With more and more companies embracing and planning their AI strategies, even companies who started early will find themselves looking for the right resources to get them going. Your business strategy will impact how you wet your feet with ML and advanced analytics. Ultimately, your forecasting, market strategy, and growth strategy will influence how quickly and in what manner you approach the creation, continuation, and growth of your AI programs internally. You might decide to start with your product, but you could also start with other internal applications of AI, such as using AI within HR or finance functions. If you recall our list of AI use cases from *Chapter 6*, you'd typically be using AI to address eight main types of problems:

- **Anomaly detection**: Applying AI/ML to identify rare or unusual patterns that do not conform to expected behavioral patterns (e.g., fraud detection)
- **Clustering/peer grouping**: Applying AI/ML to group data points into distinct sets based on similarities (e.g., user segmentation)
- **Classification**: Applying AI/ML to assign labels to data from an established set of categories (e.g., spam detection)
- **Regression**: Applying AI/ML to predict a numeric value based on training data (e.g., predicting housing prices)
- **Recommendations**: Applying AI/ML to suggest relevant items or content based on established preferences (e.g., shopping cart recommendations)
- **Ranking**: Ordering or prioritizing items based on criteria (e.g., search engine results)
- **Optimization**: Finding the best solution from a set of potential solutions (e.g., driving route optimization)
- **Generation**: Creating new data/examples based on input examples/prompts (e.g., image generators)

All of these use cases can be applied to your product or your internal processes. Feel free to skip ahead toward the end of this chapter to see our case study example, which demonstrates how regression, optimization, recommendation, and anomaly detection can potentially be implemented into new AI features for traditional software products looking to integrate AI. Perhaps you want to predict how much demand there will be for some products you sell for inventory purposes. Perhaps you want to see whether some of your customers' purchases are fraudulent. Perhaps you're trying to group your customers into targeted categories. All these problems can be helped with AI. Similarly, you could choose to use certain algorithms in your product development to help your product make predictions for your customers, show them how they are performing compared to their peers, or offer them insight into anomalies within their performance. AI is applicable to both internal and customer-facing expressions. Here's something to think about – which of these use cases is relevant to you and which approach would you choose?

> The number of AI start-ups, the amount of annual investment from venture capitalist firms, job opening projections, and the annual number of AI patents, as well as published AI papers, are all on upward trends. According to Crunchbase, 35% of all start-up investment went to AI companies in 2024, the highest percentage on record. Forecasts are showing a 28.46% increase in the AI market by 2030, where it will reach $826.7B worldwide. Currently, there are 70,717 AI start-ups worldwide. This means we are experiencing a massive explosion in demand and investment in AI, and this is still very much the beginning.

Working with an AI consultant

If you're just wetting your feet with AI, one area to start with would be to solicit the help of a consultant. When doing your own research within the space, it can be daunting to understand the spread of options you have before you. These options can span everything from data preparation to model deployment and AI program optimization. If you have limited domain knowledge in this area, working with a consultant can help you navigate the choices you have. You'll also have to make some decisions about what kind of infrastructure you want to have, what the reporting structure will look like for your high-tech projects, and where the easiest applications of data or AI projects would be for maximum visibility within your organization. I always recommend companies start small with one high-impact project because your first attempt at AI will likely be met with hesitation or resistance.

It's something unknown, and there isn't an established baseline yet for what success looks like. Consultants are great with a first-time project, particularly if the company struggles with aligning AI initiatives to the organization's goals and resources, or just getting you set up with the right infrastructure and workflow to make it happen. Many companies will attempt to take on building an AI pipeline on their own without the help of outside consultants and strategists, but if you don't have a good baseline of understanding for how to manage an AI program, you could wind up spending a lot of money on decisions that haven't been properly thought through.

Here are some tasks that could be handled by a consultant:

- Identifying AI use cases
- Developing an AI roadmap that aligns with the company goals and mission
- Data assessment and inventory
- Data preparation and best practices for data collection, cleaning, and annnotation
- Implementing data governance strategies to manage data access, privacy, and compliance
- Infrastructure, tool, and platform recommendations for building the AI tech stack
- Setting up and configuring AI infrastructure

Working with a third party

Perhaps you want to go with a consulting firm to help you instead of one individual person. This would be similar to what we've just discussed in the previous section but with the following key advantages:

- You'd have the added help of a consulting team or a firm that's done similar work in the past, has worked with a number of companies, and has a good reputation.
- Third parties also tend to verticalize, so you could search within an ecosystem of consulting firms that specialize in your specific industry or use case. Understanding your own goals in relation to your industry peers is a way to make sure you're going in the right direction from a competitive perspective.
- Working with a group of consultants also allows you to learn from a group of people who will then work with and educate your own employees.
- Many companies also struggle to find the right pool of candidates who actually have the experience they need. AI is a wide field, and depending on the specialization of your product, it might be hard to have enough support internally with the talent you already have. Consultancies have their own network of AI talent, so building a partnership with a firm you trust means you can have a third party that can guide you as your AI/ML needs change over time. It's a way to hedge against making infrastructure or business investments that turn out to be the wrong choice later on.
- It's also a way to make sure you have access to a talent pool that actually suits your needs.

One word of caution for using third parties is that it's most recommended for high-level organizational education or really select proofs of concept. If you use a consultant for your product or some internal function and you think they're going to create an awesome ML product for you and let your engineers take it from there, you're going to be disappointed and your ML project will likely fail to get corporate or leadership sponsorship because the results will be abysmal. Perhaps you're not ready to have someone on staff yet.

Perhaps you have great developers who can tune, retrain, and deploy models regularly but aren't up to the task of establishing an algorithm or choosing a model and you need someone with more of a statistical background to create the best model for your use case. You might be tempted to just use a consultancy to get you to level one. Avoid this temptation. If you're actively looking to start with data strategy and AI, use a consultant to educate yourself and arrive at a decision on how to move forward, and then invest in having someone who actually knows your business from the inside out take it from there.

The first hire

This brings us to the first hire. Your first hire will be someone who's enough of a generalist to lay out all the options for you much like a consultant would, but they're working for you full time so they can run with the ideas if their leadership team agrees with the recommendations. As you're vetting and interviewing this person, make it clear to them in which area of the business you're looking to include AI, automation, advanced analytics, or ML. I would also make it clear what the goals are for that position, what you honestly need help with, and what stage you're at. If you're in a position to only hire one person to get you started, I would also make sure you have agreement on what a first project looks like so that no one's surprised on the other end.

You should have a pretty clear idea of your data and AI strategy so that you're not overloading this person with unrealistic pressure. AI PMs, ML engineers, data scientists, data engineers, and data analysts all specialize in their own ways and it's essentially impossible to find a *data generalist*. Someone who can tell you the best infrastructure for your data and workflows, actually build those pipelines, clean all that data, load it into models, train those models appropriately, deploy those models, and finally communicate the success of all that work in a way that's meaningful to the business is impossible to find. You also don't want to just hire a whole team right out the gate without having ways to meaningfully apply their work. These roles don't come cheap, so investing in a team before you really know how you're going to use them is like buying a submarine before you get your license to drive a car. AI is a massive investment in your business, so make sure you educate yourself and take advantage of consultants before you start hiring so that you're not wasting your money. This is an ecosystem, and there's a place for everyone in it, with good reason.

The first AI team

Now let's say you do have a good idea of where to start, you have some good early projects for your data and AI folks to work on, and you're already seeing some success in your AI applications. Excitement is buzzing within your business, and you've got a number of departments reaching out to you about starting their own ML projects. Suddenly, your lone data scientist or ML engineer is overloaded with work. You start to invest in a team and essentially an entirely new department. In *Chapter 6*, we discussed PMs roles of ML engineers, data scientists, data engineers, and data analysts – this is your AI team. This is the team that will allow you to responsibly build out your AI function and optimize your infrastructure, data management, workflows, data pipelines, modeling, training, deploying, and communicating with the business. Having a dedicated team for this is optimal and ethical because a lot of work and maintenance goes into AI management. If you can afford it, you know how to best make use of your employees, and you know how to keep them happy and dedicated, they will create a wealth of opportunity and success for your product and business.

Fear is not the answer – there is more to gain than lose (or spend)

Believing in and dreaming of success are vital skills for a an AI PM. So much of the role surrounds concepts such as building a product vision, mission, and strategy and ultimately using these tools to create a roadmap that will manifest these more nebulous concepts. As a PM, you have to train yourself to visualize.

You can't visualize if you don't maintain clarity and focus on your goal. It's all about alignment. You might find yourself saying, "Do we have alignment?" over and over again as a PM to the point where you'll find TikTok videos joking about PMs saying this way too much. Perhaps what makes it so funny is that alignment is so crucial to the role that the job function itself can be distilled into this one word. You're creating alignment in all ways. You're aligning leadership, marketing, sales, customer success, operations, finance, engineering, and countless other impacted business functions around the product and helping build something successful that all these teams can be proud of.

This is the emotional part of the role, and it's one I feel the need to bring up in this section because many AI PMs might find themselves facing the cumbersome task of grappling with AI when they might not know much about it in the first place. The fear that can arise from this is the very essence of why I decided to write this book. In my first role as an AI PM, I felt fear pop up for me. I felt insecure about my own knowledge of AI in the first place. I had a background in data science and ML before my first official role as a ML and I still had this fear and uncertainty. I began managing a book club about data science, AI, and ML books to keep this fear at bay. After each book, 38 and counting, I find myself becoming more and more confident about the subject matter. I then started writing and sharing my own articles as a way to manage the complicated emotions that came up for me as I was managing these products.

When I look at these emotions objectively, I think it's was funny that I still have reservations about my knowledge and skill set. After all, there aren't too many PMs out there who have come from the data science and ML space in the first place. Because of this, depending on their level of confidence, I would venture to say that most AI PMs will likely find themselves grappling with many of the emotions I have felt. Particularly in the previous sections, which focus on non-AI-native products, I wanted to address the emotional side of AI product management and give further insight into how to contextualize the power of AI and ML to keep the fears at bay.

If we zoom out of the emotional aspect of the role and take a bird's-eye view of the jobs to be done, what does alignment really look like? We discussed many of the first strategic hires in the last section and it was tempting to include a list of tasks these roles might fulfill. But in the end, any of these roles could do any number of these tasks. A lot of fear can arise when we venture into the unknown and get inundated with conflicting messages about what the right way forward is. I feel it's beneficial instead to think about what levels of alignment you'll be working on as an AI PM, which will keep your fears at bay. Here is a list of areas of alignment you'll be working on from the beginning:

- **Planning and roadmapping**: In the early days, you'll be identifying areas where AI can add the most value to your product or organization. You'll work with your technical counterparts on finding areas to enhance operational efficiency, enhance customer experiences, or enable new business. This will be formalized in the form of a structured plan for AI adoption; it will include steps and milestones for success.
- **Data strategy and management**: You'll be working with people who will help you assess the quality, availability, and structure of your data and give you guidance on how to fill any gaps you have and collect, clean, transform, and access your data to make sure it will be effective enough for any AI/ML modeling you will do.

- **Technology and infrastructure**: You'll receive guidance on the best combination of tools, frameworks, and platforms to support the goals you've planned and roadmapped. An inventory of your technology stack will show you which areas you'll need to supplementarily fill. This will include everything from compute needs to storage and pipelines to support the flow of data to feed your ML pipelines in deployment. Managing all these vendors and partnerships will be a big part of creating a trusted ecosystem of solutions that align with your needs.
- **Model development**: You'll work with your technical counterparts in AI/ML to develop initial prototypes, proofs of concept, and MVPs in a way that aligns with the roadmap you've built, the data ecosystem you've gathered, and the technology stack you've agreed to. Model selection, learning paradigms, training, hyperparameter optimization, and performance metrics standards will all be involved in these discussions.
- **Model deployment**: You'll work with your technical counterparts in AI/ML and engineering to deploy AI models into production environments once those models are developed and tested, which means you'll need alignment on integrating them into an ecosystem that already exists. This also means you'll need to set up maintenance infrastructure to make sure those model deployments are working smoothly once they are up and running.
- **Financial management**: You'll communicate with a number of stakeholders on the financial feasibility of your AI endeavors. Gathering alignment on cost estimates and ROI projections for your AI implementations, demonstrating financial impact, and defining performance indicators to measure the success of your AI initiatives against the goals and product strategy you define early on are important.
- **Compliance and legal management**: You'll work with people to identify which regulatory/legal frameworks and compliance standards your AI/ML program will have to adhere to for your data and AI/ML model use. This will include best practices and strategies you'll use to make sure your AI/ML models and training data are avoiding bias, discrimination, and unfair downstream effects.
- **AI program management**: As your needs and complexities grow, you'll be setting up and managing a cross-functional, collaborative team of stakeholders. There will be a number of projects involved, at the very least for all the areas we covered in the previous bullets. You can think about change management and training alignment here as well, particularly if there are AI knowledge gaps in your team or if you're helping to evangelize AI skills and competencies within your organization.

Remember that these major areas all take time to align. Depending on how complex your organization is, some of these things can take years. But no matter how long it takes, anyone can help you with alignment with any of these areas. It could be a private consultant, a third-party consultancy, a first hire, or a fully fledged AI team that helps you with this. Who you choose to bring on board and create a budget for will be up to you. But by and large, these are the major areas you'll be receiving guidance and strategy on if you want to control alignment from the beginning.

Even the way we see roles defined now will begin to change. At this point in time, you might find roles such as AI PM or ML PM, which articulate the focus on an AI or ML product. Over time, my prediction is that this qualifier will start to go away as all PM roles mature with AI adoption. As you've seen time and again throughout this book, as with the idea that all companies will become AI companies by the end of the decade, the same will hold true for PMs. It will likely be that within 10 years, you'll have to have some comfort and familiarity with AI in order to be a PM at all.

The last subject I wanted to touch on in this section is the cost associated with having an AI program in place for you to research, develop, test, and manage your AI activities. As an AI PM, you're not necessarily keeping up with the costs and decisions of how to handle this function in your company, but you'll need to be familiar with the costs that go into your product. If you're actively working on developing AI features, this will be one area you're keeping track of as you build products. Part of your work as a PM is comparing the costs of certain feature developments with the potential gains from investing in those features. The metrics you choose to demonstrate the value and **return on investment** (**ROI**) from the features you're investing in will make your business case.

Remember that it's still relatively early in the AI game. Outside of the top tech companies, AI adoption is modest. This means that in this phase, and at least for the next few years, the companies that are coming out with AI features for their products are still ahead of the curve, which means it's arguably the best time for AI feature releases. Coming to market with AI features for your product means you're able to strive for more market share faster, before your competitors get the same idea and send out the same marketing message: *"We're the only AI-powered X"* will only last so long. Across many industries, AI features themselves are the marketing differentiator. Maintaining alignment and controlling our part in the areas covered above will allow us to come to market confidently. Beyond the marketing, AI features, if implemented properly, should give you a more robust, smarter product. Actual product functionality and results will broaden your product's reputation and lead to organic growth. If you're leveraging AI features in a way that helps your product save your customers money, save them time, or make them money faster, then it should be learning from what works. The inherent promise of AI/ML lies in the following equation: data plus models equals more data and better models. Your performance will improve if your data is clean and if your models are updated regularly. Getting more and more familiar with this process will fine-tune your intuition when it comes to your product's functionality and performance in the market you're in. As you begin to see progress in the scale and insights your AI features are adding to your product, you'll begin to anticipate potential issues that can arise with an AI product. This is where some healthy fear is appropriate in your AI PM role.

Anticipating potential risks

Part of the specialization that takes place in the current form of an AI PM is knowledge of what can go wrong. Your AI features may be taking certain liberties with your customers' data, and it's your job as a PM to anticipate the potential risks your product might create for your users. As a PM, you're looking to address the most important business problems your users have with your product, but you're also looking to minimize potential adverse effects.

I know, we got you all jazzed up about being a fearless AI product manager only to then bring up the things you should be afraid of because we want to make a distinction. I'd rather you be afraid of the potentially *harmful* side of AI than be afraid of the *complexity* of AI systems. Understanding the structure and complexity is the easy part. Making sure this complexity is handled properly is the hard part. For example, rather than worrying about where to start with implementing AI features, I would rather you worry about those features having potentially negative effects on your users and customers downstream. This could be related to the choices your AI features provide your users, decisions they make on behalf of your customers, or conclusions they arrive at with little explainability.

The inherent responsibility that comes with being a PM who manages an AI product lies primarily with the amazing impact AI can have. Design, optimization, and cost management are crucial for AI products. Making sure your product is learning with data samples that are representative and free of bias is easier said than done. Making sure your product is fair and inclusive of all types of users and that your data integrity is healthy is not a straightforward task. You have to also consider downstream effects that would impact your users and customers both immediately and long after interacting with your product. For example, if your product is a dating site, are you able to say the choices you are offering your users are representative of the population of users you have and free of bias?

Here are some questions to ask yourself to ensure you are a responsible steward of your AI product:

- What happens later on to your users or their data?
- Could they be harmed months or years after your AI systems have interacted with them?
- Is your product handling your users' data or their privacy appropriately?
- Are you keeping your users' data safe against attacks?
- Are there checks and balances in place to make sure there's organizational AI accountability?
- Is a human overseeing some parts of your product development life cycle, or do you only respond to potential issues or breaches reactively?
- Are your models explainable and transparent?
- Do you and your engineers have a reasonable way to account for the insights and decisions your product is making on behalf of your users?
- Can you say with good conscience that you're consistently evaluating your models' decision-making and performance?
- Are you accounting for data drift and model decay?

Depending on how advanced your product organization is, you might have AI PMs who are focused on individual products as product owners. You might have PMs who focus on the infrastructure and developer tools side. You could be an AI PM who is more focused on the research side to find new innovations to add to your product line. You could also have a role that is most focused on the responsible building of AI and the ethics of your product. If you are the only PM, you'll have to consider all the aspects mentioned above and know a bit about everything.

AI shouldn't be a silver bullet. Slapping on an AI tag just for the sake of a competitive marketing edge is not a strategy for success because if AI is not significantly improving your product and is only included as a vanity feature, you'll eventually find yourself outperformed by products that do leverage AI appropriately. Understanding which AI features to expand on and how they're improving your product will make your job as an AI PM easier and you'll understand this better and better with time. Use your curiosity and experimentation and take intelligent risks. Getting clear on how you can best support your company and what the limitations and benefits of AI/ML are will set you up for success in prioritizing features and planning product strategy. Understanding the potential harm your AI investments might contribute downstream will allow you to build credibility both internally within your company and externally within your industry.

As an AI PM, you are evangelizing AI for your organization. A big part of this indirect influence is actually creating awareness and understanding the growth, the risks, and the inherent opportunity AI offers not only to your product and business but to other stakeholders you work with regularly as well. All roles will be impacted by AI with time. You're just ahead of the curve, in a sense, because your role is most intimately connected with AI and its potential and you can use this firsthand knowledge to guide your peers through this transformation as well.

How LLMs are evolving and the rise of open source LLM capabilities

With the rise in generative AI and **large language models** (**LLMs**), we're getting close to a mass market adoption of AI in a way we haven't seen in prior years. OpenAI was founded in 2015 and it took them almost 7 years of research and development before they released the first version of ChatGPT. But LLMs have been in the works for some time. Prior to the LLMs we know of today, there were many early language models by other names:

- **N-gram models** have been able to predict the next word in a sentence since the early 2000s.
- **Feed-forward neural networks** have worked as probabilistic language models since 2003.
- Two of the models we covered in *Chapter 1*, **recurrent neural networks** (**RNNs**) and **long short-term memory** (**LSTM**) models, have been able to handle longer sequences of prompts and responses since the mid-2010s.
- **Transformer** models came out in 2017.
- Transformers laid much of the groundwork for enabling larger pre-trained models like Google's **BERT** and OpenAI's **GPT** series that came out in 2018.

Now that these pre-trained models have gotten us to the point where we can see how far we can take the tech, the emerging trends are more focused on specializing these capabilities and maximizing efficiency. We've already seen products come into the market combining text, image, audio, and video capabilities, but we're also seeing models that are able to match the performance of larger models, but with a fraction of the parameters or compute LLMs need.

For many companies, one easy way for them to incorporate LLMs into their business, online experience, or offerings is to offer a chatbot. This has led to the practice of companies bolstering pre-trained models with knowledge bases and specialized domains and tasks. Adapting LLMs and improving their fine-tuning will take us further as their capabilities are harnessed in more focused ways.

We don't yet fully know how this adoption will impact markets. For some, the threat of generative AI models, including LLMs, is existential. Others see LLMs as a tool they can use to make their jobs easier. We know it will support workers across a variety of jobs, and that it will make workers who use generative models more productive. But we still don't know to what end.

Many companies today are using the hype of LLMs as a way to take agency away from workers, in some cases downsizing. Some are encouraging their employees to use generative models while others are penalizing their use. We're still very much at the onset of generative adoption but if we take a look at the numbers, adoption is staying strong. According to an article by iOPEX, "67% of organizations have incorporated LLMs into their workflows and are utilizing GenAI capabilities to unlock unprecedented insights, automate tasks, and optimize processes." Another company, called Amperly, conducted an LLM survey and found that 37.3% of participants use AI tools like ChatGPT for work every day. Almost half, about 46%, said they use them a few times a week.

LLMs are still making their way through the hype cycle, and even OpenAI struggles to keep users coming back after the initial excitement wears off. We don't know what sustained use will look like, but we do know that tools like this are primed for open source communities to flourish and find innovative ways of pushing LLM capabilities further. Communities like Hugging Face, EleutherAI, BigScience, Turing-NLG, Cohere.ai, Stanford NLP Group, and the OpenAI GPT API community are all encouraging new concentrations of innovation to happen in their own way.

These are spaces where repositories of pre-trained models, collections of datasets for training and evaluating models, discussion forums, GitHub repositories, research projects, Discord servers, API access resources, workshops, events, documentation, educational resources, and libraries can all be found for AI practitioners to play with. As we inch closer to the next frontier of LLMs, many of the next innovations are likely to come from **open source communities**. So far, they've achieved the following:

- They've helped democratize these technologies so that we can advance faster than we have before.
- They've made a variety of models and resources listed above more widely accessible to more people and they've shared the burden of cost by providing access to pre-trained models, something that would have been hard to access before.
- They've also allowed for fast prototyping, so if you have an idea you can validate that idea with others quickly and get feedback to iterate on the next version.
- Open source communities have also allowed for diverse uses of LLMs. Because of the nature of models being developed in various languages or domains, they've allowed for more creative uses of LLMs to emerge.
- Open source communities shine a light on innovation because there are simply many eyes on it and their work is published publicly. Because many open source projects adhere to and promote AI practices that are ethical, they're able to influence the broader AI landscape for good.

The important thing to remember is the availability of these tools won't go away any time soon. Bringing more transparency and ethics to this space will allow for future use cases that are brought to market through community oversight and collaboration. If knowledgeable communities working on generative AI innovations can come together to help audit, identify, and address biases, we can explore future AI potential more safely. Investing in open source communities with your time, knowledge, energy, and ideas makes it so that the next generation of technologists will have a broader set of use cases and applications to work with as they dream up new possibilities with AI. We're also likely to see new start-ups be nurtured through open source communities, competing with the larger monoliths that have a chokehold on AI innovation today. If the future of LLMs is open and diverse, this will yield great things for the market.

Case study

In this case study, we will take a look at a fictional green tech company called GreenCo123, which specializes in optimizing energy consumption in commercial buildings. They're working on developing the next generation of their energy management system, GridOS. This latest version will include AI features. The goal of this evolving product is to help businesses reduce energy costs, lower carbon emissions, and improve building sustainability by predicting energy demand and automating HVAC systems, lighting, and other energy-intensive operations.

GreenCo123 looked into the following features for bolstering GridOS with AI capabilities:

- **Energy demand forecasting:** Time-series analysis and predictive modeling were done using LSTMs to handle time-dependent data used for predicting the energy demand of a building based on historical usage data, occupancy patterns, and operations schedules. This internal data was appended with external data like weather forecasts.
- **Automated energy control**: Reinforcement learning models were used to dynamically control HVAC systems, lighting, and other devices to optimize energy consumption across all occupancy levels.
- **Anomaly detection:** Unsupervised learning models were used to detect anomalies in energy usage to help find equipment malfunctions or trigger unusual consumption patterns, which enabled proactive maintenance.

Implementation

Initially, GreenCo123 partnered with a green tech consultancy to help them understand their options for how they can make GridOS significantly more robust to be able to compete in their new AI-evolved marketplace. From there, they also worked with this green tech consultancy to define the system requirements they would need to enable the AI features they identified together. They also collaboratively developed the system architecture, the data pipelines, and the ML models that would initially be used to power the features above with their staff data scientist, along with other stakeholders at the company.

Once GridOS was up and running and deployed on a few pilot buildings, GreenCo123 onboarded an in-house AI team to support the AI system they created and to further develop and support the needs of GridOS. They already had a data scientist with domain expertise in green energy on staff to sign off on the models used for each of the features and to work with the consulting company, so handoff was relatively seamless. But once deployment on the pilot locations started, they went forward with their team-building strategy, which happened in tiers:

- **Tier 1:** Post-pilot, there was a core team to handle the day-to-day operations of GridOS composed of the data scientist they had on staff, a data engineer to maintain the data pipelines, and an engineer to maintain the pilot program of GridOS. They also hired an AI PM to handle all these activities and bring focus and intention to GridOS.
- **Tier 2:** As they were scaling, the team grew the data science and ML engineering headcount to handle the scale of the models and of the deployments across multiple buildings and locations.
- **Tier 3:** In the third tier, they had individual teams to support activities for each of the main features discussed earlier. One team was focused on predictive modeling for energy demand forecasting, another on the automated control reinforcement learning models, and a third on anomaly detection activities.

Risks

Integrating AI was a calculated risk for GreenCo123, which had to set aside funding for the green tech consultancy as well as for the data engineering, preparation, and infrastructure development needed to support the AI features they wanted to implement in GridOS. Here are the risks that could have potentially impacted the next gen of GridOS and how the AI PM handled this added complexity:

- **System reliability and safety**: The automated energy control feature posed an issue for GridOS. What if it malfunctioned or risked exposing occupants to suboptimal conditions? The AI PM had to get ahead of that by not only implementing testing procedures that were rigorous and thorough but also introducing a manual fallback mechanism so that the AI system could be overridden easily in the event of an issue.
- **Data quality and standardization**: Because so much of the data needed to run GridOS came from different buildings that had data that had varying degrees of data uniformity, the AI PM had to make sure there was a robust data governance framework in place. This allowed the data flowing from multiple sources to be standardized and cleaned before being introduced to models for training.
- **Adaptability**: Each of the buildings GridOS was managing had its very own energy usage patterns, signatures, and equipment. The AI PM worked ahead of these issues by working with the ML engineers to prioritize the development of adaptable models that could be easily customized for each building.

Markers of success

The initial goal of GridOS to help businesses reduce energy costs, lower carbon emissions, and improve building sustainability was linked to specific KPIs and metrics within the company to gauge the success of the new AI features being added:

- One of the top metrics of GridOS was the percent reduction in energy consumption, which grew between 20% and 30% for most of their clients under the new AI features.
- Another metric that brought significant success to the next-gen platform was the occupant comfort index, which is measured by occupant feedback, which grew by 28%.

This meant the new platform was notably responsive to the needs of occupants, which also resulted in more reviews and public case studies. Because the AI PM was able to align the success of the AI feature adoption to the metrics and KPIs GridOS was already closely monitoring before the features, they were able to indicate how much of the product's success came from those features by mapping the metrics and KPIs to the timeline post-pilot.

Summary

The AI revolution is happening at many levels, and in this chapter, we took a look at a few of the major areas of how AI is impacting industries as a whole, companies from the inside out, and the role of PMs as well. For the companies finding themselves in industries that are now seeing AI transformation, the first part of this chapter focused on the various areas of AI adoption across industries and how this is affecting the future of work itself. The second part of this chapter focused on how AI is transforming companies themselves and how you can get started at the organizational level to prepare for AI adoption. The third part of this chapter then took these concepts down a level to the PM-level view and the mindset needed from the product's organization to ensure AI is adopted within a product in a way that ensures integrity and strength moving forward.

In the next chapter, we will be exploring the various ways we're seeing AI trending across industries, based on prominent and respected research organizations, in an effort to inspire PMs out there to begin to formulate their strategies for elevating their products into AI products and the considerations they must keep in mind when attempting to approach AI.

References

- *Realizing 2030: Dell Technologies research explores the next era of human-machine partnerships*: `https://www.dell.com/en-us/dt/corporate/newsroom/realizing-2030-dell-technologies-research-explores-the-next-era-of-human-machine-partnerships.htm`
- *Keep your AI projects on track*: `https://hbr.org/2023/11/keep-your-ai-projects-on-track`
- *Artificial Intelligence Index Report 2024, Chapter 1*: `https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_AI-Index-Report-2024_Chapter1.pdf`

- *AI gobbled a record share of startup funding this year*: `https://news.crunchbase.com/ai/record-share-startup-funding-2024-xai-anthropic/`
- *The future is now: AI startup statistics in 2024*: `https://edgedelta.com/company/blog/ai-startup-statistics`
- *The growing adoption of LLMs in production in the enterprise*: `https://www.iopex.ai/blogs/enterprise-generative-ai-the-growing-adoption-of-llms-in-production/`
- *LLM survey 2024: Generative AI adoption statistics at work*: `https://amperly.com/llm-survey-generative-ai-adoption-statistics/`

13

Trends and Insights Across Industry

In the previous chapter, we looked at how the rising tide of **artificial intelligence** (**AI**) is affecting companies across the spectrum, as well as the ways AI is affecting companies within their own operations. In this chapter, we will look at the various ways we're seeing AI trending across industries, through the lens of prominent and respected research organizations, in an effort to inspire product managers and entrepreneurs out there to begin to formulate their strategies by elevating their products into AI products. We will look at the key growth areas for AI integration based on the conglomeration of research and trend analysis from Forrester, Gartner, and McKinsey. In addition, we will go over the various considerations they must keep in mind when attempting to approach AI, including AI readiness and enablement.

Analyzing trends and understanding the growth areas for AI and **machine learning** (**ML**) adoption can open us up to powerful future opportunities. The nature of what we build and how we work is changing because of this massive shift in adoption. For most of this book, we've focused on AI products but we've also alluded to the various ways companies will be changing from an operational perspective due to AI. The best advice for companies looking to capitalize on building AI programs to support the products they make is to also use those same programs to leverage AI for internal purposes. If you're investing in this technology, you'll also want to use it to make internal efficiencies and elevate the performance of your product.

If you're in a position to be ready to take your product to the next level by introducing AI/ML capabilities, it means you've accepted the responsibility and privilege of supporting such an ambitious endeavor. In this chapter, we'll be taking a look at what reputable consulting, research, and advisory companies have to say about where the greater opportunities regarding AI adoption lie, as well as looking into the trends coming up in adopting AI in traditional software products. Getting a sense of what's been showing promise and the kinds of AI adoption that have been substantially improving products is helpful when looking for inspiration for how AI can improve your own existing products.

In this chapter, we will cover the following topics:

- Highest growth areas for AI integration
- Low-hanging fruit – quickest wins for AI enablement
- Riding the GenAI wave

Highest growth areas for AI integration

First, we will take a look at some of the growth areas of AI from some of the most prominent research and consulting groups. Understanding what the signals are saying gives us the motivation and foresight to be able to anticipate some of the greatest opportunities that lie ahead. This is particularly helpful in the context of revolutionizing a business or product toward AI because many product managers and technologists may be at odds about which specific areas of a product or service they might want to begin bolstering with AI capabilities. According to Statista, the global AI market grew beyond $184 billion in 2024, up $50 billion from 2023, and is expected to surpass $800 billion by 2030 by conservative measures. More generous projections from a report by Research & Markets see the market reaching over $1.1 trillion by 2030, with the greatest gains coming from the North America and Asia-Pacific markets.

Gartner's AI adoption curve identifies several stages that organizations go through when they're implementing AI:

- **The planning stage**: Where organizations identify viable AI use cases and experiment. Examples can include healthcare startups or retailers looking to leverage AI in pilot programs and use cases before they commit to a larger-scale implementation.
- **The experimentation stage**: Where companies will test AI in controlled environments to assess impact. Regional financial institutions, banks, and logistics companies may test AI tools on a smaller scale before expanding them into broader operations.
- **The expansion stage**: Where they're moving beyond pilots and scaling AI use cases. Larger retail chains and telecommunications companies may scale after initial pilots and smaller rollouts show promise and ROI.
- **The optimization/transformation stage**: Where organizations fully integrate AI to leverage it as a competitive advantage. Tech giants, major car companies, and larger financial institutions already have AI integrated deeply into their products and operations and demonstrate mature AI adoption.

As of this book's release, Gartner is projecting that by 2027, 24% of organizations globally will still be in the planning stage of AI adoption while 36% of organizations will be in the experimentation stage. Driving this adoption is the rise in AI-bolstered software products, particularly the rise of software companies expanding their AI features and capabilities, as well as a rise in AI tools made for creating other AI products and applications, and the rise in AI native products themselves. One of the general themes to have come out of research done by Forrester, Gartner, and McKinsey in the last year has surrounded the experimental nature of organizations in the last few years when it comes to AI transformation. Across the board, organizations are reporting a shift in the tides. If the last few years were about understanding use cases and experimenting with new models, applications, and methodologies, the next few will be about applying those lessons learned to drive real impact and growth.

Their retrospective *The state of AI in 2022*—and a *half decade in review* survey offers us a chance to stop and reflect on how the last five years of AI adoption have gone. AI adoption has more than doubled since 2017 across respondents, and the average number of AI capabilities that organizations use has also doubled. These are encouraging numbers that indicate adoption is moving steadily along with global trends and predictions.

The McKinsey review survey also goes into some of the top use cases of AI deployment, listing that service operations optimizations (internal applications of AI in organizations) have been the highest functional activity among respondents who have adopted AI, followed closely by the creation of new AI-based products. Other notable functions included were new AI-based enhancements of products, product feature optimization, and predictive services and intervention.

Some of the top growth areas we will look at in this section are:

- **Embedded AI**: AI that's applied and integrated into the operations of organizations and foundations of products
- **Ethical AI**: Considerations of responsibility and privacy in AI deployment
- **Generative AI (GenAI)**: Generative and Web3 use cases of AI
- **Autonomous development**: The growing field of AI-generated code

Let's look at each of these areas in detail.

Applied/embedded AI – applied and integrated use cases

Applied/embedded AI relates to AI being integrated into core operational activities within a business. Despite all the success and hype of other uses of AI in recent years, demand for embedded uses of traditional ML applications remains high. Whether it's for predictive modeling, optimization, or numerical tasks, applied/embedded AI is performing the bulk of productivity-enhancing tasks across industries and is one of the highest growth areas of AI.

If your product is geared toward B2B internal processes such as architecture, operations, fulfillment, supply chain, HR, or procurement, you might be interested in bolstering your product with AI features that help companies fulfill their commitments to their consumers. Based on Forrester's research on 2025 predictions, AI platform budgets are expected to triple, "*driving a 36% compound annual growth rate from 2023 to 2030*." We will continue to see more B2B use cases and applications in embedded/applied AI use cases, but the more your product can offer real value to your B2B consumers, the more you will be able to help enterprises shrink the latency between insights, decisions, and results. This area of AI, which Forrester refers to as *AI Inside*, deals with the integration of AI that's embedded into operations.

In B2C contexts, having applied AI bolstered by GenAI means organizations can transform their marketing and customer experience to revive customer loyalty. By using applied/embedded AI in contexts that anticipate the needs of your customers and help them in their own decision-making, you can deliver value through applied/embedded AI use cases. According to the Forrester *Predictions 2025* report, limiting the data gap between marketing and loyalty practices will help B2C companies make the most of applied AI. Today, most B2C companies keep these areas of their data siloed, but you won't be able to make the most of personalization and customer insights without marrying the two and leveraging AI with the data.

Often, the most value your product will be able to give your customers is actionable information they can use to make high-level decisions. If your product is dependable in delivering this, it is worth the hassle for companies to invest in their own AI infrastructure to achieve the same results. Gartner outlines this concept in their *applications-centric AI* framework, where they consider innovations in engineering, decision intelligence, and operational AI systems.

For Gartner, the focus of this area of AI is toward helping improve decision intelligence overall in order to both reduce the technical debt and visibility for businesses internally and help reduce the risk and unpredictability of outcomes. In other words, the most impactful part of this area of AI, according to Gartner, is that it helps decision makers have both clarity of their internal processes and make their most important decisions using insights that are derived from their own data.

Gartner predicts a rise of **causal AI**, which are systems that identify and act on *"cause-and-effect relationships to go beyond correlation-based predictive models and toward AI systems that can prescribe actions more effectively and act more autonomously."* This report also includes things such as the **digital immune system**, a set of practices and technologies that focus on optimized resilience in the face of threats that would impact an organization's ability to bounce back from any threat to their ability to deliver a customer or user experience. As they see it, this would cover a wide variety of tools that help businesses prepare for potential risks, which learn from past failures to better prepare for the next. Rather than looking at this from a cybersecurity perspective, this would pertain to any internal process that could negatively impact experience and delivery to a customer. The **applied observability** category also highlights tools that look for any deviations within an organization's operations that would impact core business functions and infrastructure.

McKinsey featured *Applied AI* as one of their top strategic features in their *Technology Trends Outlook 2022* report, which they define as the intelligent application of AI toward solving classification, control, and prediction problems in order to automate, add, or augment business use cases. Some of these use cases are defined in their report as risk management, service operations optimization, and product development. Based on their research, the factors contributing to this include the global expansion of AI adoption, more affordable AI implementation routes, an improvement in training speed, high growth in innovation based on patents filed, and a marked increase in private investment growth for AI-related companies.

McKinsey has isolated the areas of AI that are showing the most noteworthy promise for AI adoption:

1. They list the general umbrella of ML at the top of the list for its use in optimization problems and its ability to learn from data using statistical models.
2. **Computer vision** is listed as the second most noteworthy area for using visual data for facial recognition and biometrics.
3. Third is **natural language processing** (**NLP**) and its prevalence in speech recognition and virtual voice assistants.
4. Fourth is **deep reinforcement learning**, specifically for things such as robotics and manufacturing lines.
5. Lastly, they list **knowledge graphs** for their ability to derive insights from network analysis.

Now let's take a closer look at what the data shows us. McKinsey reports that 56% of their survey respondents said their organizations are adopting AI, there's been a 94.4% improvement in training speed since 2018, there were 30 times more AI patents filed in 2021 than in 2015, and $93.5 billion was invested in AI-related companies in 2021, doubling from 2020. McKinsey also cites industrializing ML, or the adoption and scaling of ML capabilities, as another key trend that will show the greatest return for the next decade.

They also outline industrializing AI into several key areas, including data management, model development, model deployment, live-model operations, and overall ML workflow. Based on their research, the global revenue impact potential of AI across industries is $10–15 trillion, and companies adopting AI are 2.5x more likely to experience higher 5-year returns to their shareholders. These growth projections do come with a massive caveat, however: 72% of the organizations they surveyed for this report have not been able to successfully adopt and scale AI. Among the biggest contributors to this are difficulties transitioning pilot projects into products, models failing in production, difficulties with scaling their AI/ML team's productivity, and limitations in containing risks.

McKinsey's 2024 *Technology Trends Outlook* reinforces the trend by saying "*applied AI and industrializing machine learning, boosted by the widening interest in gen AI, have seen the most significant uptick in innovation, reflected in the surge in publications and patents from 2022 to 2023*." We've covered most of these major areas of ML extensively in this book, and it's no surprise to see them repeated over the years.

Ethical AI – responsibility and privacy

Another high growth area for Forrester is the niche of **responsible AI** or **ethical AI**. Thanks to advocacy groups and increasing pressure on lawmakers to define ethical uses of AI moving forward, the area of responsible AI is seeing a boom. AI/ML capabilities will go through increasing waves of regulation because of how pervasive AI has already become in virtually every area of human life, even in these early days. This opens up a lot of opportunities for products and services that already deal with issues of fairness, bias, and governance to then move into the area of governing AI bias.

According to Forrester, the more AI adoption we experience, the more *"existing machine-learning vendors will acquire specialized responsible AI vendors for bias detection, interpretability, and model lineage capabilities."* This is a rosy outlook for companies already operating in this space, as we're likely to see demand for these kinds of services increase as the next decade moves forward.

Considering our category of embedded and applied AI is the top growth area for AI adoption, we will increasingly experience ML embedded into more and more of the products we use, whether we're individual consumers or businesses. With the increase in adoption will come an increase in how heavily the use and deployment of ML are scrutinized. Gartner agrees, with their Research VP highlighting the growing need for AI that's reliable, understandable, impartial, and accountable. Responsible AI helps address biases in data, fosters trust through explainability, and ensures regulatory adherence, even in the face of AI's inherent unpredictability.

Among these general comments about AI transparency and fairness, we also see a slew of categories of prediction, including the rise of dynamic risk governance tools, cybersecurity mesh architecture, and cloud sustainability, as well as the rise in decentralized identities and the move toward consumers owning more and more of their digital identity and data as Web3 applications take off.

The numbers for the growth of ethical AI practices are also encouraging and indicate that individuals and companies are taking the ethics and responsibility that AI comes with seriously as their level of experience and investment grow. Forrester predicts 25% of tech executives will need to report their AI governance activities to their boards to cover "*explainability, fairness audits of high-impact algorithmic decision-making, and accounting for environmental impacts of AI (green AI)*." This also indicates that boards and high-level decision makers are aware of the risks that come with not properly accounting for the potential harm AI can pose to customers, users, and organizations.

Based on their own research in this area, Gartner had predicted that by 2023, anyone hired for development and training work on AI systems would need to demonstrate their expertise in ethical and responsible AI. Unfortunately, it remains to be seen if this is the case. Their outlook on what they consider **causal AI**, or AI that's able to demonstrate explainability and deliver on communicating causal relationships, is much longer. Gartner believes causal AI could take up to a decade to achieve mainstream adoption but will eventually have a "*transformational impact on business*." Here, we also see the relationship between trust and responsibility going hand in hand, with companies understanding that explainability and communicating risk impacts even those that hold high positions within companies.

Gartner lists **AI trust, risk, and security management** (**AI TRiSM**) as a strategic technology trend for 2023, citing that in the US., UK, and Germany, 41% of organizations had experienced an AI privacy breach or security incident. This further highlights that the world is still ripe for responsible AI practices that bleed into privacy and security and they are very much needed and will continue to stay relevant. Seeing these indicators in research organizations is encouraging, but it's just a prediction of the future, not a manifestation.

As companies continue to build and invest in their AI products, they will need to not only build products in a trustworthy way for the benefit of their users but they should come to understand that not doing so can threaten their own investment in the AI tech they've spent so much time, money, and patience building.

McKinsey's *Technology Trends Outlook 2022* report also highlights **trustworthiness** and **explainability** in embedded/applied AI as one of the top considerations when incorporating AI into products. The ability to hedge risks, particularly as AI use cases expand, is of particular concern. Adhering to laws, remaining vigilant on ethics, and building AI with social robustness are their recommendations for building in a way that mitigates harm. They also note that explainability is closely related to this, further highlighting that as models become increasingly complex and are deployed on high-risk use cases, the issue of explainability will remain critical.

McKinsey lists explainability as threefold:

- Explaining how the model actually works
- Organizations demonstrating causal explainability, which is about explaining why certain outputs come from the inputs.
- Trust-inducing explainability, i.e., explaining why you can trust and deploy a model, which is perhaps the most important

Building AI/ML products and upgrading traditional software products into AI products will be a continuous process of reinforcing their trustworthiness, reliability, and safety across multiple areas. It gives me confidence that these three prominent advisory organizations give ethical AI the importance it deserves as we continue to build the future with these powerful tools.

GenAI – immersive applications

Another theme surrounds the rapid adoption and success of GenAI. Acquiring valuable lessons and putting them into practice with the power of GenAI causes pressure to prove ROI and demonstrate tangible value for organizations. I want to encourage you as AI PMs to follow suit in planning bold roadmaps for your products and avoiding distractions by prioritizing use cases that ultimately deliver real value to your customers, users, and internal stakeholders. But keep in mind that GenAI models have been democratized for only a few years. According to Forrester's Q2 2024 *AI Pulse Survey*, "*49% of US generative AI decision makers said their organization expects ROI on AI investments within one to three years and 44% said within three to five years*."

Because of the evolving nature of the AI technology landscape, long development cycles, and the resource intensity of building AI products, expecting true ROI within three or five years is still ambitious. AI PMs need to help their organizations balance the pressure to prove ROI with innovative, bold roadmaps that balance opportunities for short- and long-term ROI. Building a strategy around re-investing gains from features and capabilities that bring short-term ROI into longer-term projects is a great way to build on AI successes.

McKinsey's *Technology Trends Outlook 2024* report highlights GenAI and the context windows for LLMs, which grew from 100,000 to two million tokens. This means that as generative models are being adopted at record numbers, their ability to handle more and more complex prompts has scaled massively. Now we're seeing LLMs grow from text summarization and image generation to more sophisticated capabilities through video, images, audio, and text. This means that as the capabilities of powerful GenAI models grow, they're being integrated more aggressively into software tools and customer-facing applications.

Generative AI apps had a great year from 2022 to 2023. Adoption was strong and the general public flocked to generative tools that flooded the market in efforts to take the new tech for a spin. Curiosity and experimentation were strong. But as we discovered from the open trials mentioned earlier, spurred on by the SAG-AFTRA strikes in 2023, 2024 was a more sobering year for GenAI. There are many more alongside the open legal battles mentioned. The list is quite long today. Most well-known GenAI companies are currently being sued, and for most of them, the issues lie with the creativity the models are training on themselves. Public perception of GenAI and its uses will continue to fluctuate as we grapple with the implications of having these powerful models have access to a multitude of human creativity.

If they're training on human creativity and the humans who taught them everything they know are cut out of the equation, what does that mean for the future of the feedback loop? Eventually, there will be so much AI-generated content out there that the models will be starved of innovative, new, human creativity to learn from once again. In and out of the court, some agreement will have to be reached. We can't have human creativity and the totality of human artists and creatives against GenAI models. There will have to be a symbiotic way for both to exist and be celebrated for what they bring to the world in their own way. Forrester predicts there will be an increase in Fortune 500 companies' dependence on GenAI and how they will generate content with AI tools because "*human-produced content creation will never be fast enough to address the need for personalized content at scale, and in the next year, we expect to see at least 10% of companies invest in AI-supported digital content creation*."

This indicates that organizations will increasingly rely on these apps for their own marketing and content strategy goals, signaling opportunity in the B2B marketing world for GenAI apps. But this number doesn't include the many consumers who will be interested in GenAI apps for purely recreational purposes. Gartner's data follows suit, citing that by 2025, GenAI will account for 10% of all data produced, up from less than 1% today. Gartner's focus on GenAI can be used for a range of activities, such as creating software code, facilitating drug development, and targeted marketing, but it can also be misused for scams, fraud, political disinformation, forged identities, and more. This further highlights the increasing dependency businesses will have on using GenAI, not just from a marketing perspective but also from an operational perspective.

The more business use cases are built and optimized for GenAI apps, the more content they will be able to generate, particularly when you consider those needs starting to exist at scale. Those numbers are hard to predict for the consumer market, however, because 2022 was really the first year GenAI saw huge levels of adoption, with individual consumers using it for recreation. According to Gartner, the primary focus of GenAI initiatives is 38% customer experience and retention, 26% revenue growth, 17% cost optimization, and 7% business continuity. If the current buzz around GenAI apps in the consumer market is any indication, we will likely see more excitement around new use cases and products, particularly as they integrate with other emerging technologies such as AR/VR, Web3, and metaverse applications.

Autonomous AI development – TuringBots

Low- and no-code tools have been around for some time now in the hopes of bridging the divide between the technical needs of enterprises and the cost and time needed to onboard a dedicated staff of developers. This is now being pushed a step further with the emergence of what Forrester calls **TuringBots**, or bots that write code. Gartner predicts this will be a major theme in the coming decade and refers to these as **machine learning code generation tools**, but the idea here is the same. These are ML models that would work alongside human developers and are integrated into development environments in order to offer suggestions on the code base using descriptions in natural language or fragments of code as a prompt. Because of the highly adaptable nature of ML to optimize repetitive tasks, it's unsurprising that ML is leveraged to optimize the very way it's built.

Gartner goes on to predict a rise in what it calls **adaptive AI** or systems that *"continuously retrain models and learn within runtime and development environments based on new data to adapt quickly to changes in real-world circumstances that were not foreseen or available during initial development."* This means that AI that's more sensitive to real-time changes in the training data is also able to adjust its parameters and goals in an effort to make them more autonomously nimble. According to Gartner's 2024 technology trends report, GenAI will "*significantly alter 70% of the design and development effort for new web applications and mobile apps*" by 2026.

McKinsey's *Technology Trends Outlook 2024* report highlights next-generation software development as another strategic trend to look out for, citing the increase of low-code/no-code platforms and AI code recommendations based on context from natural language, AI-based testing that automates performance testing, and AI-assisted code review as another key trend for the next decade. They see the strategic involvement of AI in the planning and analysis, architecture design, development, coding, testing, deployment, and maintenance phases as a key focus for organizations that are building products or internalizing this trend to experience its benefits.

The data for autonomous AI development is also encouraging and mirrors the data we're seeing for GenAI apps. Because of the rise in popularity of reinforcement learning and language models such as GPT-4 (which powers ChatGPT), we're seeing that we can take low- and no-code tools a step further and actually produce code with the help of AI using nothing more than basic instructions expressed in any natural human language. Forrester's 2024 *Developer Survey* states that 24% of executives plan to leverage GenAI across the entire software development lifecycle.

McKinsey follows suit, citing that 70% of new software development will be built using no- or low-code tech by 2025, which will reduce development time by 90%. This is great news for the build and creation process of software. The deployment time will be twice as fast because of continuous integration and continuous delivery practices, which will be helped by AI. Considering deployment is often the hardest part of making ML and AI use cases successful and also brings a lot of confidence. 37% of McKinsey's survey participants said they would use AI/ML to test and maintain their existing code repositories, bringing the advantage of AI-assisted coding to another level.

So far, we've had a chance to discuss the highest growth areas, embedded AI, ethical AI, GenAI, and autonomous AI development, coming from some of our most highly reputable research and consulting organizations. In the next section, let's get into the numbers to get a better sense of the scale of these trends.

Low-hanging fruit – quickest wins for AI enablement

By now, you've seen how involved applied ML is for an organization to embrace fully, and we've just spent most of this chapter looking through the various growth areas in AI for organizations that are already in business and are looking to capitalize on these growth areas. In *Part 1* of this book, we discussed the various layers of infrastructure that need to be supported in an AI program. In *Part 2* of this book, we discussed the AI-native product. In this current part of the book, we're discussing the transition of incorporating AI into a traditional software product. This means we can now move on to set the stage for what this adoption looks like.

Before we get started, a caveat: we really can't talk about AI transformation unless we also set ourselves up for success to be able to begin the long and arduous process that is AI adoption. We have to make sure the conditions are right. Whether you're working in an organization that's ready to build an AI-native product or you're in a traditional software environment that's ready to adopt AI at the feature level, there's a level of expectation setting that's required for the endeavor to be successful. But setting expectations properly in an already established software company that isn't used to the demands AI will place on an organization is especially difficult.

There are both tangible and intangible challenges to building the right conditions for AI. The tangible ones lie with the infrastructure, the investment, and, frankly, the skill. Building out your team and bringing in folks who can troubleshoot and navigate the tricky conditions AI will force you to work in requires a special set of skills, and those skills are sought by virtually everyone right now. Building out established processes and workflows that support the various facets of an AI program is a formidable project that can be challenging to scope, much like AI projects themselves.

The intangible challenges are more cerebral and emotional. In *Chapter 7*, we discussed the differences between traditional software products and AI-native products. If you recall that section, one of the biggest differences lies in the uncertainty AI brings. When you have an established team that's already used to a certain level of expectation, it can be difficult to articulate just how much AI is going to change things. The uncertainty that comes with AI makes it very difficult to communicate the full extent of allocating a budget for teams, purchasing software, managing expectations of timelines and costs, and resetting boundaries in terms of how teams work with one another.

You can't have a fully supported AI program if your internal teams don't have the right infrastructure and procedures in place. You also can't have it if you do have the infrastructure and processes but your internal teams haven't emotionally internalized them yet. According to Gartner, *"By 2025, the 10% of enterprises that establish AI engineering best practices will generate at least three times more value from their AI efforts than the 90% of enterprises that do not."* There's a significant upside for companies that do their due diligence to set up their AI programs properly.

The work of bridging the gap between the tangible and intangible challenges of AI is what's referred to as AI enablement. In order for a traditional software product to transition into an AI product as seamlessly as possible, its leaders and product managers need to create a plan to make this transition successful. Adopting AI is a massive shift for your organization and the way it builds products. It's a fundamental change to how we expect to build. Building a strong culture of AI enablement means that you're preparing your teams as best as you can for what's to come. Let's look at some companies that have been able to do this successfully:

- **Adobe** embedded Adobe Sensei into their Creative Cloud, Document Cloud, and Experience Cloud product suites by focusing on data quality, infrastructure, and managing expectations internally of how AI would transform workflows.

- **Domino's Pizza** was strategic about its AI rollout by using AI to optimize operations as well as customer experience through the DOM Assist chatbot and predictive ordering algorithms. They were able to be most effective by developing clean data pipelines, adopting continuous data transformation, and aligning organizational expectations around AI integration and launching Project 3TEN, an initiative that aimed to have pizza ready for delivery after three minutes and delivered to the customer within ten minutes.
- **Capital One** uses AI in its mobile app experience and in their fraud detection systems. Their tangible challenges related to building the necessary data pipelines and supporting infrastructure to support their AI endeavors, as well as training the team on working with AI and even fostering trust in AI decision-making. They even went as far as to launch skill development programs to equip their PMs with the right skills to effectively support AI deployment and become model owners in collaboration with the data scientists they were working with.

These examples highlight that companies are able to be most successful with AI integration when they prioritize their data pipelines and infrastructure needs, and when they prepare their teams culturally through setting expectations, training, workshops, or other upskilling opportunities.

At its core, AI enablement is really all about the data. Finding a way to most efficiently collect, label, and curate large amounts of unstructured data is what allows model performance to improve most consistently. Improving that data pipeline, keeping it clean, and making sure there is a steady supply of it to be used for training purposes is what gives the AI/ML pipeline its capabilities. Making sure the right data transformations are happening and the proper level of data quality is in place is the ultimate objective of AI enablement. This is because of the tremendous reliance AI has on data.

The best combination of AI enablement or AI readiness will come from having strong use cases to support your investment in AI, having strong networks of clean data you can leverage to support those use cases, and having a clear governance strategy of who can access what and how roles are delineated within your organization to demonstrate ownership, control, and safety measures to make sure your AI investment is as successful as possible. We will go into these concepts in more depth in the next chapter.

Riding the GenAI wave

Now let's turn our attention to GenAI models and tools. The rise of LLMs and generative models has helped democratize access to AI, bringing capabilities to companies of all sizes without the significant investment they originally needed to scale their products. It's not the models themselves that are becoming more accessible, but the infrastructure needed to employ those models in production as well. Given the significant costs to train and deploy applications of AI in production, **infrastructure as a service** (**IaaS**) platforms exist today to help drive innovation and adoption. Some of its uses are:

- **Serving as a knowledge base**: Since popular LLMs like ChatGPT have come out, we've seen them being deployed as chatbots and virtual assistants across a variety of products and company websites. In many cases, they serve as a centralized knowledge base. Though a generalized approach might be helpful for expanding support hours or coverage, the companies that are innovating with chatbots the most are able to map chatbot activity to specific use cases that reduce or limit reliance on call centers and support teams internally.

- **Increasing efficiency through automation**: Incorporating generative models toward the end of automating routine or repetitive tasks allows employees to focus on activities that are more strategic and contribute more to the company's established goals. The key here isn't to try and do it all at once. It's to choose a few use cases that create the most bottlenecks for the company, or for specific teams, and to monitor how performance from generative models actually improves those bottlenecks. Operational efficiency is one of the main ways companies will make use of generative models in practice.
- **Achieving high levels of personalization**: AI can be even more effective if we're able to build personalization into the chatbot experience to the point that they're able to understand their users enough that they can offer recommendations that are relevant and helpful to them. Personalization doesn't have to be limited specifically to recommendations. It can be expanded to include marketing materials that are tailored to specific groups of customers and users. In some cases, this level of personalization can be achieved at the individual user level based on their behavior patterns and established preferences.
- **Product development**: We also see applications of generative models on the product development side. In many cases, they can be used for brainstorming, understanding market analysis, market trends, testing, and design. When combining generative models with IaaS platforms, product teams can get to work on ideating and vetting AI features, prototyping and testing new concepts with key customers before rolling them out in a way that's even more cost-effective. This reduces the amount of money and time product teams invest in their product development lifecycles to bring features and capabilities that may make their products more competitive in their markets faster. Some of the key aspects where generative models can be used are:
    - **Managing costs**: IaaS services are often on pay structures that are cumulative, so you're paying more as you use their GPUs, TPUs, and data storage. They also allow for more flexible budgeting across the seasonality of their product. If there are periods where they have higher customer demand, they'll pay more, and vice versa. This means that product teams can have more control over their time and resources as they relate to building AI features and capabilities than they otherwise would. Previously, a company might have had to invest a lot of time and energy into building out an AI cost center or center of excellence in order to start experimenting with AI. Today, with the accessibility of generative models, LLMs, and IaaS environments, they can build more comfortably and see what works well early.
    - **Scaling:** IaaS environments also help significantly with scaling. The technical debt of building a code base that's flexible enough as you scale customers 10x or 100x isn't ever going to be completely gone, but you'll have more control when catching up. What this means is companies can focus on scaling without the concern that their product experience will suffer as they make efficiencies. This allows companies to potentially achieve a global audience earlier. In the case of continuity issues, particularly as you scale, IaaS providers offer backups, disaster recovery, and data availability when something does go wrong, reducing the impact felt by your customers and users.

When leveraging generative models with IaaS platforms, you're able to focus your AI program or product team on elements that will impact your product development lifecycle the most. Upgrades, compute, resource allocation, and infrastructure management can be handled by your provider so that your product and dev teams can focus on building and shipping products. Because many companies will be building with this mix of products and services, that means product teams will have to be focused squarely on bringing valuable capabilities to market. Their time is better served spent on product innovation to stay competitive.

Let's review a couple of successful examples of GenAI integration:

- **Salesforce** took the approach of integrating GenAI through its Einstein AI platform, which uses ML models to offer enhanced recommendations, predictive analytics, and customer insights within its CRM suite of products. It leveraged IaaS platforms like AWS to scale its AI models efficiently without having to worry about those costs upfront. This meant that it was able to offer personalized customer experiences, optimized sales processes, and improved operational efficiency.
- **L'Oreal** took the approach of integrating GenAI through its product development cycle by incorporating it into its ideation and trend analysis workflows. It also used GenAI models to personalize beauty product recommendations for its users and customers through a skincare and cosmetics device called Perso. It also uses IaaS services to manage the large datasets that power these tasks, as well as model training and experimentation. It's been able to bring new products to market faster through its work with GenAI models and their enhanced recommendations have improved sales and customer experience.

In these examples, we see a focus on using GenAI models and IaaS services to deliver value at a much faster rate than otherwise would have been possible with more traditional approaches. By strategically using AI models within scalable cloud environments, organizations can minimize the burdens that can come with optimizing AI adoption with less flexible infrastructure. AI PMs will always be weighing the benefits and the costs of going with one solution over the other. When adopting GenAI in their product development lifecycles, AI PMs can ask some of these questions:

- How will the cost structure of infrastructure or IaaS impact our long-term AI roadmap?
- Will any significant releases or milestones need to be altered to prioritize the budget?
- How can we optimize resource allocation to minimize overspending during peak demand times?
- How can we best set performance monitoring thresholds to trigger automatic scaling and manage compute costs?
- How will our infrastructure/IaaS scale?
- Will sudden surges in usage be normalized through autoscaling mechanisms?
- Will there be latency related to AI models or data availability issues as demand surges?

Summary

This chapter was all about trends and insights for AI adoption collected from some of the most reputable companies that speak about it. We looked at some of their insights and projections for the coming years and decades regarding AI adoption. Building an AI-native product is, in many ways, more straightforward than transitioning a product from traditional software development to an AI product.

In this chapter, we wanted to set the stage and discuss some of the high-growth areas for AI adoption because, for many companies, knowing where to begin is often the hardest part. Once you're in the flow of things, you can better anticipate what comes next, but when you're at the precipice of a major paradigm shift, there's a lot of friction. Going over the growth areas, data, and common use cases and setting the stage for AI enablement was an intuitive choice to make sure AI PMs out there are aware of what adopting AI can mean for their product. It's a big responsibility and undertaking, and it should be treated as such.

In the next chapter, we will go into a lot of the applied considerations of transitioning a product to adopt AI and tangibly what that means for your product team or organization.

References

- *Predictions 2022*: *Successfully riding the next wave of AI*: `https://www.forrester.com/blogs/predictions-for-2022-successfully-riding-the-next-wave-of-ai/`
- *What's new in artificial intelligence from the 2022 Gartner Hype Cycle*: `https://www.gartner.com/en/articles/what-s-new-in-artificial-intelligence-from-the-2022-gartner-hype-cycle`
- *What's new in the 2022 Gartner hype cycle for emerging technologies*: `https://www.gartner.com/en/articles/what-s-new-in-the-2022-gartner-hype-cycle-for-emerging-technologies`
- *McKinsey Technology Trends Outlook 2022*: `https://www.mckinsey.com/~/media/mckinsey/business%20functions/mckinsey%20digital/our%20insights/the%20top%20trends%20in%20tech%202022/mckinsey-tech-trends-outlook-2022-full-report.pdf`
- *Predictions 2023: AI will become an indispensable, trusted enterprise coworker*: `https://www.forrester.com/blogs/predictions-2023-ai`
- *Gartner identifies four trends driving near-term artificial intelligence innovation*: `https://www.gartner.com/en/newsroom/press-releases/2021-09-07-gartner-identifies-four-trends-driving-near-term-artificial-intelligence-innovation`
- *AI is a new kind of creative partner*: `https://www.forrester.com/blogs/ai-is-a-new-kind-of-creative-partner/`
- *In a world first, South Africa grants patent to an artificial intelligence system*: `https://theconversation.com/in-a-world-first-south-africa-grants-patent-to-an-artificial-intelligence-system-165623`
- *Gartner identifies the top 10 strategic technology trends for 2023*: `https://www.gartner.com/en/newsroom/press-releases/2022-10-17-gartner-identifies-the-top-10-strategic-technology-trends-for-2023`

- *AI software spend will double from 2021 to 2025, reaching $64 billion*: https://www.forrester.com/blogs/ai-software-spend-will-double-from-2021-to-2025-to-reach-64-billion-and-grow-50-faster-than-the-overall-software-market/
- *Forrester forecasts AI software will grow 50% faster than the overall software market*: https://www.prnewswire.com/news-releases/forrester-forecasts-ai-software-will-grow-50-faster-than-the-overall-software-market-301638854.html
- *5 ways Forrester predicts AI will be "indispensable" in 2023*: https://venturebeat.com/ai/5-ways-forrester-predicts-ai-will-be-indispensable-in-2023/
- *Forecast analysis: Artificial intelligence software, worldwide*: https://www.gartner.com/en/documents/4007140
- *Notes from the AI frontier: Modeling the impact of AI on the world economy*: https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy
- *The state of AI in 2022—and a half decade in review*: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review
- *McKinsey Technology Trends Outlook 2024*: https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-top-trends-in-tech
- *Gartner experts answer the top Generative AI questions for your enterprise*: https://www.gartner.com/en/topics/generative-ai
- *Gartner top 10 strategic technology trends for 2024*: https://www.gartner.com/en/articles/gartner-top-10-strategic-technology-trends-for-2024
- *Focus areas for preparing a learning and development strategy for 2024*: https://elearning.adobe.com/2024/01/focus-areas-for-preparing-a-learning-and-development-strategy-for-2024/
- *Capital One: Transforming traditional banking to an AI-first experience*: https://d3.harvard.edu/platform-digit/submission/capital-one-transforming-traditional-banking-to-an-ai-first-experience/
- *Domino's Pizza Enterprises delivers in record time using AWS for predictive ordering*: https://aws.amazon.com/solutions/case-studies/dominos-case-study1/
- *Salesforce announces Einstein GPT, the world's first Generative AI for CRM*: https://www.salesforce.com/news/press-releases/2023/03/07/einstein-generative-ai/
- *Unveil Perso, the world's first AI-powered device for skincare and cosmetics*: https://www.loreal.com/en/news/research-innovation/unveil-perso-the-worlds-first-aipowered-device-for-skincare-and-cosmetics/

14

Evolving Products into AI Products

In this chapter, we will explore the areas that benefit and negatively impact AI adoption in existing products. Every product will be different and every product owner will come to their own conclusions about what level of AI adoption to infuse into their product. For some products, something as simple as adding one AI feature is enough. For other products, a fundamental change to the underlying logic that powers your product might be required. The decisions of how to transform your product should, first and foremost, be determined by your product strategy and should serve your overarching company vision. These decisions should be collaborative and should come with a high level of executive sponsorship.

As the chapter continues, we will explore various areas of AI transformation, and they will serve as a step-by-step guide to building a product strategy that supports the evolution of your product. We will look at how to approach brainstorming about your AI options, which considerations will be most important as you brainstorm, how to evaluate your product's performance within the competitive landscape, how to build a product strategy that sets you up for success, and, ultimately, which signposts and milestones will help you build confidence and which will indicate that you need to go back to the drawing board.

We will discuss these in the context of the stages of AI product development we introduced in *Chapter 6*: ideation, data management, research and development, and deployment. Whether you're adding AI features or upgrading the existing logic of your product, having an established plan that supports your product strategy will be the best way to successfully update your product for commercial success with AI in a way that's sustainable. Once you've achieved the successful transformation of your product into an AI product, feel free to return to *Part 2* of this book where we go over the various aspects of building and maintaining an AI-native product.

In this chapter, we will cover the following topics:

- Ideation – what's possible and what's probable
- Data management – the bloodstream of the company

- Competition – love your enemies
- Product strategy – building a blueprint that works for everyone
- Red flags and green flags – what to look and watch out for

Ideation – what's possible, what's desirable, and what's probable

There are multiple ways in which a product could stand to benefit from AI adoption. Understanding the Venn diagram of what's probable, what's desirable, and what's possible is an important part of your AI product strategy. Clarify the distinction between the ideal outcomes you'd like to achieve (desirable), the technical feasibility of those outcomes (possible), and the viable options you can explore based on current constraints, data, and priorities (probable).

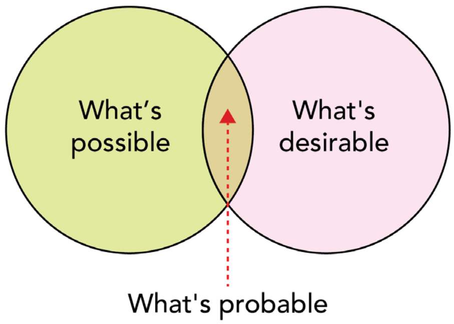

Figure 14.1: The Venn diagram for ideation

As an AI PM, you must asses your wish list (what's desirable) against the constraints you face (what's possible) to identify the probabilistic or most realistic choices that you can then focus on. As you continue this exercise, you'll go through a spectrum of possibilities. You'll start with a really open-ended, right-brain brainstorming session on the cerebral notion of value and how your product can deliver it to your customers, which will then be refined by a left-brain, analytical breakdown of the fruits of that brainstorming session. Both are important parts of the product management creative process. More specifically, this exercise represents the *Ideation* phase of the AI product development cycle.

Whether you look at your product from the perspective of the main problems your customers face, the *jobs to be done*, or feature parity, you'll want to get a sense of which AI enhancements would be most high-value for your product, which would be the most low-cost, which have the best data readiness or availability, and which have the highest degree of executive or market sponsorship. Since this section is about how to evolve an existing product into an AI product, take your time with this step because, with it, you're coming to an important decision: identifying which problems or opportunities AI could offer significant value to. Being mindful about taking your time to really familiarize yourself with the models, the benefits, and the options you have before you will be important.

Before we get into all that, a word of caution about adopting the term AI. We've discussed the notion of starting with your *why* when it comes to AI and how important it is. We've also discussed the notion of the marketing buzz that comes with AI and how attractive it is to investors and customers alike. In the current AI climate, there's a lot of noise that's generated when AI and ML enter the lexicon of your product's communication efforts. You don't want to fall victim to the trap of attracting a lot of attention without necessarily having the substance to support that attention. A big part of that is getting clarity on the feasibility of using AI for an identified problem or opportunity.

The last thing we want is for PMs and technologists to go through the effort of investing in AI and push out vanity features that don't actually make their products better only because they want to capitalize on the attention AI offers them. Examining the availability of data, infrastructure needed, and potential technical limitations means you're being honest about what the proposed AI features can even do for your product. For that reason, we recommend grasping the real opportunities that AI and ML will afford your product, assessing technical aptitude and ability, and then creating a series of lists that will rank the order of these potential opportunities according to three key categories: *value*, *scope*, and *reach.*

The ultimate goal of brainstorming and creating these lists is to define a potential path to AI maturity within your product. Understanding where your current product falls short, which areas you want to start with bolstering with AI, and which areas you want to transform as part of your overall product strategy and roadmap over time will be an important part of the process. AI transformation won't happen overnight. You will have to start somewhere and, given the cost of investing in AI, being strategic and intentional about where you want to start will make all the difference in the long run.

In the following sections, we will be applying various lenses to ideations when it comes to AI adoption. Those lenses will relate to how AI adoption will impact your product based on perceived value for your customers, how AI adoption will impact the scope of your product, and, finally, how AI adoption will spread to your end users.

List 1 – value

The work of product management often requires a high level of trade-off. You're always judging the amount of effort and skill required to bring a feature to light and comparing that with the proposed upside for making that building investment. If you do, however, start with the highest-value AI enhancements, you'll be coming from an intuitive, customer-centric place. Understanding first and foremost what makes your product valuable to your customers and your market and then finding use cases of AI to expand on that value should be the first order of business.

You'll need a source-of-truth list like this later on when decisions become more difficult and – when they inevitably come up – you maneuver political discussions with leadership and other key stakeholders at your company. The reason why your value list should be your primary source of truth is that every product fundamentally needs a **value proposition**: a statement about why your product brings value to its customers and users. The world has enough false advertisements, overcommitments, and exaggerations, particularly in tech. The 2019 *State of AI* study examined European companies claiming to have cutting-edge AI products and found that up to 40% of them were not substantially using AI . This is further echoed in a 2023 Eurostat commission, which still reflected only 8% of businesses were using AI in a meaningful way in Europe. Build a solid list of true AI enhancements that will deliver real value to your customers. Build an AI product you can be proud of. Start with what really matters.

That value list shouldn't just be a list of features you want to add to your product. It should be a summary of the types of value your product offers its users and customers. Features are descriptive and matter-of-fact. Values are closely tied to worth. If you're having trouble understanding the difference, start with the features themselves and then see whether you can go a step further with each feature and write down *why* that feature actually matters to your customer. If you can do this, then you'll be able to do this exercise when evaluating the products that yours will be competing against. Try and get a sense of whether or not the products in your competitive landscape offer the same level of value to your market. Another way of looking at value, particularly with the competitors that have actually gone to market with their own AI features or upgrades, is to understand how those decisions might have impacted value perception in the market.

Treat this as an exercise in *divergent thinking*. We spent a considerable amount of this book going over the various use cases and models that can be employed in ML and deep learning. We've also looked at a number of verticals, industries, and high-level adoption trends. How can you apply these strengths to your own product? Get creative and try not to stop logistical considerations from getting your ideas flowing. Anything that feels like it will be inherently valuable to even just a small subset of your customers can and should make this list. If it's helpful, consider the various types of models we covered in earlier chapters and what value they could bring to your product as it currently stands. Take stock of the benefits that various types of ML can potentially bring to your product and see where they might be able to improve on existing use cases that pertain to the product and your customers. When you're thinking about value and how to grow it in your product, use this list exercise as a way to get all the potential ideas out there, no matter how outlandish or unrealistic. Keep in mind the various strengths of AI: stacking, ranking, optimizing, predicting, grouping, comparing, automating, standardizing, and learning from ongoing trends and analysis. With due diligence, you should come out of this exercise feeling inspired about the potential directions your product can take based on its current value proposition. With some luck on top of that due diligence, you should also experience the new benefits AI can bring. There will be plenty of opportunities to whittle this list down and realistically refine it later. Marty Cagan's book *Inspired: How to Create Products Customers Love* outlines the "Four Big Risks" for product teams to address as they develop products. One way of whittling your list of ideas down later will be to assess them against the value risk. Cagan defines this as the risk that customers or users might not actually value the proposed feature or product enough to buy it or use it regularly.

One of the best ways to mitigate this risk is for product teams to focus on whether or not they are solving a real problem or fulfilling a genuine need with the solutions they put forward. Developing a product or functionality without validating demand for it first will result in high cost and energy output with very little to show for it. You can work on market research, prototype testing, and customer interviews to validate the ideas you generate in this stage.

For now, let the creativity flow and think about how you can build on the existing value your product already offers its customers and users. You will be rewarded for being less discerning during this step because you never know which of the fringe ideas on this list might become valuable seeds for future features. Brainstorming about potential features and use cases is also some of the most fun an AI PM can have – figure out a way to make this a standard part of your product strategy work and revisit this on a quarterly or biannual basis.

Some questions on value you should be able to answer at the end of this stage are:

- How well do I understand how my current customers use the product at a feature level?
- Which features are the most valuable and frequently used by my customers?
- What are the top customer pain points this product team wants to tackle with AI?
- What real customer problem will this proposed AI feature solve?
- What is the core value proposition of my proposed AI feature?
- How is AI adoption improving the customer experience of this feature?
- What types of AI/ML models or techniques will be most beneficial for this feature?
- What are the risks of customers not valuing this feature?

List 2 – scope

Once you have a list of the highest-value AI enhancements from the previous section, you can then start to prioritize them into a list of most to least scope. Scope includes not only the time and effort required to complete a body of work but also the cost and skill level required to do so. Given that we're talking about AI products here, it may not be as straightforward to understand what will require the most effort, time, cost, and skill. Often, we don't know this until we're well on our way to building out these features, but do your best to create enough of a stacked order from most to least effort for this list.

This list won't be an absolute source of truth, and it doesn't have to be. It also doesn't necessarily mean that the items that involve the least effort will get built first, but it's important to establish which kinds of enhancements will require more of an investment from the company to complete. Understanding the variety in terms of effort is most important from a planning and resource management perspective. If certain features need to be built, and they require a lot of effort, time, or skill that you might not conveniently have, knowing this will help with forecasting and making the necessary plans in the future to obtain them.

It will also allow your leadership team the time and focus to really internalize what this investment in AI will actually require of them, and whether or not you currently have the skill sets and capabilities to even embark on this path. Taking the time to discuss the scope and get clear on the organization's current level of ability is an essential part of this step. The act of acquiring the right tech resources to make the loftiest AI dreams a reality will likely be an arduous journey in and of itself. Your leadership team needs to be prepared for that.

This isn't just the case with acquiring talent – it's also the case with understanding what internal processes need to change. Often, when it comes to AI, managing expectations can be the most difficult part of making an AI transformation successful, so anything you can do early on to prime your leadership for the trials and tribulations that might impact you as you begin the process of leveraging AI will go a long way when hurdles inevitably come up.

As any PM out there knows, the scope is an important factor in deciding what to build and when, and you'll want to make sure you're cementing yourself as the authority where leadership is concerned when it comes to the AI features you propose to add to your product. Even if you do have a mature team of data scientists and ML engineers around you who you can trust and who have a ton of experience and AI wisdom, as an AI PM, you will still own the product and be responsible for the ultimate strategy related to your product. It's important that you maintain this high level of oversight and command over your product's key AI considerations and the major decisions impacting your product because you are the captain of that ship. Congratulations! It's a gift and a curse.

This responsibility also relates to another of Cagan's "Four Big Risks," the feasibility risk, which focuses on whether the product or capabilities can actually be built with the resources available. There will always be some combination of time, skill set, technical ability, or organizational capacity you're juggling to make any idea come to life. Being honest about addressing feasibility as early as possible means technical stakeholders will be working on something that can actually deliver on its promises.

If you can't defend the decisions made in relation to your product, you're not the authority, and this will only cause confusion and undermine your leadership later on. Remember that your technical team is there to support the product, offer sincere advice, and give their best estimates on the scope and on how to build the features that are decided on by the strategic arm of the organization. They are there to build and implement the decisions that you make in collaboration with your leadership and stakeholder team. Having a clear definition of ownership and using your technical team to accurately reflect on the scope to the best of their ability is what this exercise is all about.

Some questions about scope that you should be able to answer at the end of this stage are:

- What is the time/effort required to develop, implement, and maintain this AI feature?
- What is the expected ROI for developing this feature, and has technical leadership bought in?
- Are there missing skill sets and resources to build this AI feature?
- How will you plan to fill the gaps in skills/resources to build this AI feature?
- What internal processes need to change to support this AI feature?
- What are the main risks associated with not delivering this AI feature?

List 3 – reach

By this point, you have two lenses through which to look at your proposed product AI embellishments – the lens of ultimate value and the lens of scope: effort, time, skill level, and cost. As you may notice, with our second list, we have moved from a divergent mindset and added more analytical, convergent restraints. With this third list, you add another layer of conditions through the lens of how many of your existing customers it will impact – or your reach.

Being able to create this list will require you to have enough intimacy with your existing product as it currently stands, right down to the feature level. Understanding the baseline of how your subset of customers uses your product and which features have been most instrumental to them in their experience of your product will be necessary.

You can't accurately predict the reach of these proposed value items from your original list if you don't already understand how your customers experience your product. You also won't be able to articulate to your leadership team why certain AI embellishments will be valuable to your customers if you don't already understand how they use your product.

This seems like a really obvious statement but you'd be surprised how many PMs out there know very little about how their customers actually use their own products. AI PMs can get caught in the trap of building their product team into a feature farm, and few of them actually stop to make sure their product is being built holistically and sustainably, but maintaining intimacy with their customers and users is a big part of being a PM. If you aren't using product analytics or regularly checking in with your customers, either through direct customer interviews or through in-app/in-platform surveys, then you don't really know how your customers experience your product.

If you don't have a strong grasp of how your customers use your product and what they find most valuable about it, then you're limiting your reach and building in the dark. This darkness exposes you to two more of Cagan's "Four Big Risks": usability risk and business viability risk. Usability risk pertains to the effective use of your feature or product. Even if some customers see the value of your feature or product, they may struggle to actually make use of it. If it's not intuitive, easy to navigate, or self-evident from a user expectation perspective, you'll lose those users. That's a practical reason to lose reach. Another of Cagan's risks is more philosophical in nature. Business viability risk alienates certain users based on the company's overarching business objectives. If the product can't support a company's financial model, legal constraints, brand reputation, or long-term strategy, you could lose reach that way as well.

If you've gotten to this part of the exercise and it's just now dawning on you that you don't have this information, then you'll want to start doing your homework and start this list creation exercise all over again. This, once again, ties back into the original value items we put on our first list. If you do find yourself in this situation, there's no harm. Keep your original list. Get to know your customers and what they most love about your product and start the exercise over again. This is one of the beautiful things about product management. You can always find out more! There are deeper and deeper truths to uncover about your product the more you build and the more customers you take on. Never get discouraged. If there are unanswered questions, it's only an opportunity to learn.

Some questions you should be able to answer about the reach of your product at the end of this stage are:

- How many of my current customers will be directly impacted by this AI feature?
- Which customer segments will benefit most from the proposed AI feature?
- How will I make sure this feature is intuitive and easy to understand/adopt?
- How does this AI feature align with the business's long-term product strategy and product roadmap?
- What evidence (data, research, feedback) supports the demand for this AI feature?
- What are the potential challenges or risks in prioritizing this AI feature?
- What legal, ethical, or regulatory constraints could limit the reach of this AI feature?

Case study

In January 2024, Shopify introduced a series of AI-powered merchant and customer experience features that would help streamline operations for the online retailers they support. We will use the framework that we've just discussed to describe how Shopify identified high-value AI enhancements, assessed scope and feasibility, and understood customer reach and usability. Here is a list of AI features that Shopify incorporated for the AI transformation of their platform:

- **Merchandising features**: Ability to store more complex product catalogs (along with more variety for color, sizing, and SKU options), combine listings, and create networks. Improved classification taxonomy to reduce time/effort to create product listing with automations for standardized categories and product attributes reducing manual, repetitive work.
- **Storefront performance features**: The launch of a new web performance dashboard that shows you how your store's page ranking can be improved to sell more. They also improved infrastructure speed and expanded the global footprint of buyers.
- **Semantic search**: Gives more relevant search results that customers would be more likely to buy from, even without using specific keywords.
- **Digital advertising:** The launch of Shopify Audiences, which uses ML to improve digital advertising, find more customers, lower **customer acquisition costs** (**CACs**), and improve **return on ad spend** (**ROAS**) with custom targeting lists that can be used on ad platforms like Mega, Google, Pinterest, Snapchat and Tiktok. Smarter retargeting algorithms were also launched, which are at least doubling retargeting audience sizes.
- **Shop Campaigns:** They rebranded their prior feature of Shop Cash offers into Shop Campaigns as a risk-free customer acquisition program to bring new customers into shops without the merchants having to build new audiences or create new assets.
- **AI code of ethics**: Concerns from merchants that AI adoption was moving faster than ideal in their industry were met with Shopify providing an AI code of ethics for their merchants, and it's reported that 78% of them have adopted it.
- **Shopify Magic**: A suite of free AI-enabled tools that includes:

a. **Media Editor**: A generative AI app that lets you generate product images instantly and for free

b. **Smart Sidekick**: An AI-enabled assistant that allows you to manage productivity, improve workflows, improve decision-making, and limit time on operational tasks

Now let's break this down a bit and discuss Shopify's adoption of AI through the lens of value, scope, and reach.

Value

Shopify wanted to focus on delivering features that would help merchants (their customers) better understand their own customers and, in turn, be able to offer them better conversions for the products they're selling on the platform. Since Shopify is a B2B company, the team went about this by conducting extensive market research and customer feedback sessions to make sure they understood what could make the biggest impact on their customers' business operations. Through this process, they found that real-time personalized product recommendations could make the biggest impact on improving customer engagement and expanding sales. What is perhaps most inspiring about this approach is the AI features posed a win-win scenario. Shopify won by fitting the bill in terms of offering AI enhancements that their customers found valuable. And their customers won by using valuable features that also brought them better customer experience with their own customers and more sales.

Scope

Shopify wanted to assess the required resources for implementing the AI features they settled on and took inventory of the technical skills needed to develop and maintain the AI/ML models and capabilities they were integrating into their product. Their challenge was to integrate these new features without disrupting the user experience that already existed on the platform. The Shopify engineering team then had to work on the new features in a way that would make them easily scale and integrate while maintaining the current platform performance. Shopify also had to fill gaps in their skills capabilities by hiring additional data scientists and ML engineers to support the new AI-powered features. They also put time and resources into upskilling their existing teams through training programs focused on AI capabilities and technologies.

Reach

Because Shopify was already using advanced analytics to understand how their merchants (customers and users) were already using the platform and what challenges they faced when running their business operations or working with clients, they were able to lean on that baseline of data to understand the potential reach these new AI features and capabilities would offer their customers. It also meant they were able to empathize early on with their top concerns, which means their proposed solutions were applicable to as many customers as possible. Shopify also conducted beta testing with a smaller select group of merchants to make sure they were meeting usability standards and getting feedback from real merchants in the process. Post-launch, Shopify also used user feedback and performance metrics to continue to improve the AI features they launched.

The result and impact of this work led to improved customer engagement on their platform. Merchants reported higher conversion rates due to more effective product recommendations, and customer interactions improved through automated support. Beyond this, Shopify won leadership clout and branding support in the e-commerce space through their innovations with AI adoption, which sets them apart as a leader in their industry. Shopify's ability to integrate AI features, many of which are free, while considering scope and reach using a strategic approach enhanced the company's reputation as well as its product's value.

This case study is exemplary and aspirational for a lot of AI PMs out there thinking about how to have an impact. It's possible that, with time, they'll learn more about what makes their platform valuable to their merchants and they will start to monetize upsell and advanced features more heavily. But by coming to market with an AI-improved platform and offering these features generously, they're helping all their merchants get more comfortable with AI adoption as a whole. It's a reminder that you don't always have to be so focused on the ROI of specific features. ROI can be secondary when it means broadening the horizon for the entire ecosystem with a launch.

Now that we have our priority list of potential sources of AI adoption based on value, scope, and reach, we can move on to the next big consideration for evolving a traditional software product into an AI product: data readiness. Making sure our data is in a digestible format to be able to handle the new AI capabilities is going to be the next big focus area for PMs looking to make the jump into AI.

Data management – the bloodstream of the company

Before PMs can begin the work needed to start building and developing their product, getting it ready for testing, and releasing it to their customers, they need to get really clear on the strategy of how they will position their product. This is the thought process behind the Venn diagram exercise we saw in the previous section. Now that we've gone over the process of how AI PMs can approach potential AI embellishments, we can add one more layer of scrutiny to this list. This additional layer focuses on the **data**, which is what will power every single item on these lists, and it represents stage 2 of the AI product development cycle.

Once AI PMs begin the work of understanding what they can do with the data sources they have and which data sources they'll need to make the items on their list a reality, we're getting close to actually having a plan. This stage is about gathering the right data that will be used to train and test the AI/ML models from internal systems, third-party providers, customer interactions, or external IoT devices like sensors, processing and refining that data for AI/ML readiness, and securely storing and protecting that data.

In the following subsections, we will be addressing the key areas of data readiness:

- Preparing and researching the data you have available
- Assessing the quality of the data you have and partnering with your data team
- Using the data for benchmarking the current and future adoption of your product and defining success

These will all be necessary steps to ensure you've covered your bases when it comes to data readiness and availability for AI. Not adequately taking care of any of the below areas could result in biased training data that is not representative of real-world conditions. This means the data feeding your AI system could make it potentially inaccurate and unfair. Let's look at these areas in detail to make sure we can prevent this.

Preparation and research

Doing our due diligence with regard to research around our data sources and what insights they can deliver for us is part of the work we need to do to truly be ready to make decisions on AI integrations. Your data readiness and availability will also have a huge impact on the executive and market sponsorship of your newly emerging AI product. You might also find yourself at the point where you've done these exercises only to realize that your data is not ready for the evolution of your product. Getting to that point might be an intermediary step to being able to launch AI version 2.0 of your product. Making sure you have flows of incoming data to support the features you want to go live with is the ultimate measure of readiness.

At this point, you haven't created a product strategy for how you're launching your newly minted AI product. We're still in the early phases, so we want this phase of data discovery to come after the brainstorming exercises we discussed in the previous section. The reason for this is that you can start to think about how to best gather, clean, and organize data in the ways your AI program will require of you as a foundational step that will inform your product strategy. This is an important part of deciding on the core function and audience of your product or your product's AI features. Taking stock of what data capability you currently have and how far you need to go to get to where you want to take your product is a step everyone faces when they're scaling their existing product into an AI product.

Sure, your customers have come to know and understand your product as it currently stands, and there are some sets of established use cases already, but those use cases are likely to evolve with AI. Understanding what the current data pipelines that make your product function are and how those data pipelines need to change and mature to support your new ML pipeline is an integral part of the AI transformation process. You can't avoid this step because if you do kick that can down the road and just let problems arise later on, you won't get very far as all you're going to get out of your models is nonsense.

Some questions you should be able to answer at the end of this stage are:

- What data do you already collect? How will it be used for your potential AI system?
- What relevant data will you need to start collecting to power your new AI system?
- Is this data internal, external, or a mix of both?
- Do you have enough volumes of data to sufficiently train the AI models you're considering? If not, how will you augment your dataset?
- How will you store and retrieve all the data you will regularly need to power your AI system?
- Do you have a plan for how data will need to be restructured for your AI models?
- What data quality and cleaning changes will your AI system necessitate?

- Are data security and privacy benchmarks being met? How will those change with your AI system?
- Which data governance for data management, ownership, and access standards will you implement for your AI system?
- Is your data prepared for the next phase (R&D)?
- Are there enough features in your data to train the AI/ML models you'll be using in your AI system?

Key challenges and common issues for AI PMs in this phase surround issues of data readiness, availability, and coverage. Merging data from various sources into a digestible format for AI/ML systems, having insufficient data, and planning for scalability in the future are all concerns for AI PMs who are looking to navigate the complexities of their data ecosystem. Here is a set of common issues that AI PMs can face at this stage:

- **Data and AI goal misalignment**: Making sure the data you are using has the right features your AI system needs to perform optimally is crucial in this phase. Sometimes, the data you have might lack features that are essential for the AI use case you're trying to solve. Take inventory of the data needs that will power your AI/ML capabilities and make sure those are reflected in product requirement documents and other sources of documentation.
- **Compatibility issues in data integration**: The AI/ML capabilities for your product may often require a variety of internal and external data sources that can cause downstream complications in the data pipelines. When working with the data engineers who will enable your AI/ML capabilities, make sure that they are accounting for all the data that will be coming into your AI system and that they have cleared all the relevant data sources.
- **Anticipating data needs**: AI PMs need to accurately assess what data volumes their AI system needs to power their AI/ML product capabilities. How much data will be enough for training needs to be discussed with ML engineers and data scientists. This is also the case for forecasting future data needs. AI PMs need to make sure these important discussions about current and future data attributes and volumes are discussed with relevant stakeholders ahead of time.
- **Maintain buy-in**: AI PMs need to be aware of the risks of not managing expectations around data effectively. If your leadership team thinks the data challenges at your organization are too costly or complex to solve, it could threaten their sponsorship for your product's AI/ML capabilities. This is also the case for data restructuring needs. If data needs to be restructured to meet new requirements for your AI system, this could impact budgeting and planning for your roadmap.

Ensuring quality partnerships

It's become a cliche by now but within AI, data science, and ML circles, the age-old tagline is *garbage in, garbage out*. This is an *a priori* truth of AI. Your models won't give you the insights and value you're looking for if you don't adequately prepare and clean your data. Should you find yourself in this situation and ignore your AI/ML team's recommendations, you're going to be left defending an internal AI program that your leadership team will perceive to be a waste of time and money. And they would be right, it would be a waste of time and money.

Rather than looking at your data pipeline strategy as the necessary input for your product, in an AI product life cycle, you'll come to see it as the partner that co-creates your AI product alongside your models and developers. Your data pipeline is going to constitute a huge part of your AI/ML program and it will complicate stage 3 of the product development cycle (research and development) as you're evaluating different models if it's not implemented correctly. You will need to think of it as a program. Orchestrating and coordinating your data sources and the quality of the data will be a huge part of your success with AI. Acknowledge that adopting AI is going to constitute a paradigm shift for your organization, and if you were able to skate by with incompatible silos of data wrought with duplicates and poisoned data wells in the past, this isn't going to be the case when you adopt AI.

The biggest resistance that organizations face when it comes to data quality comes from inside: the internal resistance to keep processes as they've always been. Doing nothing is easy. Keeping things as they are is the natural instinct for most people. Even if your organization is super passionate about maturing its data practice and you have champions across all major stakeholder teams within your organization, it still won't be enough. You will need to evangelize the importance of shifting in mindset when it comes to data quality, as well as the cost and time associated with making this happen. It's one thing to make a verbal commitment to improving data quality and another to actually see that manifest tangibly in the behaviors of teams.

Expect internal resistance to this and do your best to communicate the risks of empty promises. Remember that when you're changing workflows, processes, and departmental habits of a product that's been around for a while, you have entire teams of people who have been used to doing things a certain way. Making those changes will be a collective effort. This isn't just a product evolution that will happen because leadership said so. This is a product evolution that will require partnership from all employees who touch your existing product. Take time to manage your own expectations for this as well. There might be days when you're frustrated and weeks when you're not seeing changes and you return to past bad habits. It's a process and it will take time for this to be fully internalized.

Building and maturing a data team will also take time. How much time it takes will be up to you and your leadership team. You might want to use a data strategy consultant or have in-depth sessions with your technical stakeholders to really understand how to organize your data pipelines in a way that will empower AI adoption best. You may even want to jump in straight away and establish a dedicated cross-functional data team that will be responsible for the oversight of your data. Having dedicated folks to work out the details of centralizing the various silos of data you have and deciding on the best way to centralize that data and send it to your models in development, and ultimately in production, will be a great way to work through the various inconsistencies in your data that will inevitably come up.

At one of my recent roles managing data products, the time our organization had allotted for managing data quality was impossibly short. Despite extensive communication with leadership to manage the risks of not investing more fully into managing data quality issues, we were asked to press on with the time limitations. But in the end, the dashboards and data tables were unreliable because the work was unsustainable. After assessing the results and planning, we came up with a data quality model that had our team working more closely with data stewards from various departments to be the subject matter experts for data within a certain domain.

The process still took some time – at least six months of work to get to the underlying issues for each department's data, but we were then left with data pipelines and tables that we could all rely on for subsequent derivative products. The impact of our work meant that all other products thereafter would be coming from data we could trust, data that was properly vetted. The partnership between the data team and between representatives from virtually all other teams meant that we could reshape the culture around our data organization. This experience also showed us that we sometimes have to insist on what we know is right, especially when leadership is resistant to devoting time or money to projects that are well worth the effort.

Treat this as an investment in your R&D and product design phase and give it the time it deserves because it's a foundational part of your AI infrastructure. How you store your data, how often you're calling it, how often you'll need to train your models, and how you'll be planning a release schedule to support this new AI infrastructure will all be a part of this discussion. At this stage, it's just a discussion. No decisions are being made yet, but you will want to start this discussion as early as possible because, by the time you are ready to start considering models, training, and testing, you're going to want to be sitting on clean, rich data to feed those models.

Understanding your data, exploring it, experimenting with it, wrangling it, preparing it for feature development (for ML models), using it in modeling, evaluating the performance of those models, and, ultimately, deploying it are all going to be regular discussion points. So, starting conversations about data strategy early on before any real decisions about how you're going to leverage AI are finalized is a strategic move. It also allows for technical voices to become a part of this conversation sooner than they might have expected because adopting AI is a highly technical investment and those voices should be present early on.

As we've seen in prior chapters, feature engineering is a major part of making the models you choose to employ in your product successful. This further reinforces the idea that you need to make sure your technical teams weigh in on important decisions about what kinds of data are being included in your training sets.

Some key aspects that quality collaborations with your data and team will help you address are:

- Maintaining strong data integrity practices to make sure you're not having to retrace your steps when you start the R&D phase of the AI product development cycle.
- Establishing cross-functional alignment so that all stakeholders work together to create processes, reduce internal resistance, and deliver workflows that support AI capabilities.
- Long-term stability with regard to strategy and scalability means that teams are invested in sustainable AI infrastructure so that teams can plan ahead safely.
- The importance of treating AI adoption as part of R&D means businesses will appropriately invest in the right data infrastructure, data strategy consultants, and AI cross-functional team building.
- Making sure the organization invests in building data and AI/ML literacy across departments so that teams can collaborate effectively.
- Whether data pipelines are clean, centralized, and scalable enough to support AI features.

Benchmarking and defining success

Data is already the bloodstream of a company, empowering everything from internal decision-making to product performance:

- Your customer and product data will help you understand which parts of your product are already valuable and will help you deduce which proposed features will have the most impact on your customers.
- Your historical internal data will help you understand which features have traditionally taken the longest to manifest in your product to predict future scope.
- Your product analytics will help inform which UI/UX and design changes might have the most impact.
- Your training and performance data will help you understand which ML models work best for your customers.
- Your operational data will help you understand which deployments have been most successful and why.

Data helps you manage your release schedule and product roadmap and helps you scale your product. Establishing an existing benchmark using data will be an important step in creating a narrative around your plans for adopting AI. This is because, in future quarters and years, you want to be able to demonstrate that AI has helped move the needle on some of your most compelling use cases both internally and externally. You won't be able to know how far you've come later if you don't have an established baseline from the very beginning.

As an AI PM, you'll want to speak from a position of strength when you demonstrate the milestones your product has been able to hit with AI. The business will evolve and metrics will come and go, but make the most of the data you already have before you embark on your AI adoption path so that you can truly reflect on your progress when the time comes.

You'll also use your internal data to set metrics, which we covered in *Part 2* of this book, in order to define what success looks like for your product. Defining success will be a collaborative act and it's important for AI PMs to bring in all major stakeholders of the company to align on defining that success. If the product evolution into AI doesn't include the voices of your go-to-market, sales, marketing, strategy, engineering, and leadership teams, it won't have the appropriate level of sponsorship it will need to truly be effective, funded, and evangelized internally and externally.

Perhaps most importantly, without the right level of sponsorship, the success of your product will also not be tied to your most important business objectives at the organizational level. Building the right level of data fluency with these teams and using data to support your claims about the direction your product is taking is non-negotiable. Be patient with yourself, your organization, and your technical talent on data. Data is what powers this entire endeavor we're discussing. Don't give in to the pressure to rush things. There are always business imperatives to get to market as quickly as possible. The instinct to just start and build as you go will always be something AI PMs and leaders alike will have to push against, but don't rush this.

Here are some questions you can ask to determine whether you're geared for success:

- Do we have a clear baseline of data for comparison once we onboard AI features for things like customer behavior, product performance, and internal processes using existing data?
- Have we defined specific measurable outcomes and metrics that will indicate the success of the AI features we might want to add?
- Are all major stakeholders (go-to-market, sales, marketing, engineering, leadership) aligned on these metrics?
- Is there sufficient executive sponsorship for potential AI features?
- Is there adequate funding, commitment, and advocacy from leadership to ensure the success of your AI transformation long term?
- Does AI transformation support overarching business goals like revenue growth, operational efficiency, or improving customer satisfaction?

Competition – love your enemies

Now that we've discussed the internal process of ideating potential areas of AI adoption within your product, as well as the stages of ensuring your data readiness to adopt AI, we can move on to your external environment and internal development.

Your competitors will offer a landscape of possibilities when it comes to what AI adoption will look like for your product. We don't advise mimicking your competitors' AI strategy but, rather, taking what AI adoption looks like for your competition as input when you weigh up options. AI PMs who mimic the competition risk their product's failure because doing so makes them lose what makes them special to customers, investors, partners, and employees. Your product is unique and that's what attracts people to it. Own your unique value proposition and find a way to make it even more special with AI. Using the data you already have will take you quite far, but there will always be a need to append this data with a feedback loop from the outside world. Understanding your competitors will help you inform your strategy. It will give you examples from your peers who have already made the jump that you're looking to make. Some of the examples you see from your competitors will be sources of inspiration and some examples will help you avoid certain mistakes. Doing your due diligence when researching your competitors, particularly those that have also embraced AI, will have some influence on what you choose to build and should be one of the factors that influence the lists we established previously.

Some will refute this point and say that you should build a product based on your own understanding of your market and your customers' problems. The notion that you should focus on the problem you're looking to solve without the sometimes eclipsing influence of the competitive landscape is well-meaning but highly flawed. We don't exist in a vacuum. Products are compared with others in their peer group and, in many cases, one product defines another. Whether you're selling to consumers or other businesses, your product is going to be compared with similar products. Being aware of your competition doesn't necessarily mean you will build like them but it does offer you a perspective that's valuable and realistic. Ignoring the impact of your competitors' influence on your product would be like ignoring gravity. It's going to impact you whether you think it will or not.

Another important point it would be remiss not to mention here is to recognize that a big part of your competition includes the past iterations of your own product. You're competing with how you've done things. You're competing with past versions of your product. Thinking critically about this past version, particularly the version that doesn't leverage AI, is going to be a big part of your design process. Think about what the limitations are, and what opportunities, strengths, and challenges the old version of your product offered your customers. How can you capitalize on those opportunities, limitations, and challenges with this new version you're building? Treat the old or current iteration of your product as its own threat to your current product.

As you think about your competitive landscape, you might be tempted to get down to the feature level and do a comparison matrix with your top competitors. This is often a worthwhile exercise, particularly for your sales, marketing, and engagement teams later down the line, but at this stage, it's a bit too soon for that. Try to analyze your competition from the lens of the original value list we started out with.

Getting to know your competitors is a great way of not only understanding possible directions you can take with your own product but also getting a glimpse into where the industry that you work in is headed. Understanding the competitive landscape you're playing in and how it's changing can have a strong impact on your ability to see trends within your own industry and decide whether or not you want to be an active participant or an agent of change for those trends. Going to market is an entire specialization in and of itself, and as an AI PM, you're wearing that hat to quite a large degree. The better acquainted you can get with the market you're serving and how that market is evolving with AI, the stronger position you'll be in when it comes to making the call on how to leverage AI in your own product.

So far, we have brainstormed potential areas in which AI can be adopted in your existing product, grasped data readiness, and evaluated what AI adoption looks like in your competitive landscape. These are all valuable inputs you're going to consider when it comes to crafting your product strategy now that it's ready for its AI makeover. In the following section, we'll cover how to build that new product strategy that aligns with all the ideating we've done.

Product strategy – building a blueprint that works for everyone

By now, we've brainstormed potential ideas. We've taken an inventory of our data, as well as the insights that can come from our competition and greater market, and we're finally able to come to the drawing board and build a product strategy that will reflect the next major era of our product. Going from a traditional software product to an AI product is no small feat and it should be treated as a massive overhaul of the product because so much of how we build, what we build, and how we store, collect, and use data will be majorly renovated.

Building a product strategy will directly influence your product roadmap, and this will help you realize which parts of your product will need to transform first to succeed in the actualization of your product's AI transformation and commercial success. However, there's a big difference between creating a **product strategy** that fits in with your product's new value system and creating a **product roadmap** that reflects that strategy.

The roadmap is a proposed timeline of major milestones and features that will be built in the year ahead. The roadmap should not come before your product strategy is agreed upon by your key stakeholders. The roadmap is derived from this strategy and not the other way around.

We don't build what we can and hope for the best. We agree on what to build using the collective knowledge and wisdom we've gathered about our product and how to best go to market with AI capabilities. In the following subsections, we will cover the product strategy process and a proposed list of how to develop a product strategy that works for all your major stakeholders.

In the section that follows, we will cover product vision, product strategy, and product goals and how they all relate to each other. Then we will explore how to best execute a product roadmap that supports the growth of AI features and capabilities. This isn't necessarily a how-to because the order and the contributions you take in will depend on your specific product. Rather, this can be viewed as more of a checklist to ensure you're headed in the right direction:

- Product vision
- Product strategy
- Product goals
- Product roadmap

Product vision

We won't go into the specifics of what having a company mission statement or vision is about, but suffice it to say that there will be one. If it doesn't exist, this is a great moment to generate a simple statement that encapsulates why your organization exists in the first place. What is it trying to achieve and what is the vision? From this point, your product vision is about aligning with your organization's higher-level goals for the market it's serving. In your product vision, you're looking to simply explain how you see your product serving its target market. What are the needs and problems your customers and market face and how does your product meet those needs? Try to nail this down before you think about how AI shapes that vision, and then take a moment to reflect on what changes when you bring in AI. Try to whittle this down to one sentence if you can – two if you absolutely need to.

Gathering feedback and crafting a product vision are all collaborative efforts that are not done exclusively by the PM. Rather, the PM brings the right voices together to make sure a collective sentiment is agreed upon by the stakeholder team. You should be the person that organizes the meetings in which these things are decided and it's your job to make sure you leave those meetings with something everyone present can live with. You'll meet with your stakeholders regularly in those product strategy meetings we already discussed and you'll have many opportunities to reevaluate how this vision evolves. If your product vision is the "why" behind your product, your product strategy will be the "how" and your product goals will be the "what." A product vision is, well, visionary: it represents some ideal future state or ultimate impact your product will have on users or the market. It's a statement that will align everyone's efforts to that ideal future state.

Here are some examples of the product vision of well-known companies:

- **LinkedIn**: "Create economic opportunity for every member of the global workforce."

- **Slack**: "To make work life simpler, more pleasant, and more productive."
- **Ikea**: "To create a better everyday life for the many people."

For a more nuanced discussion on product vision and its impact on managing alignment with leadership and company values, feel free to refer to *Chapter 11*.

Product strategy

Knowing and agreeing on why your product exists allows you to contextualize all the activities you're doing to brainstorm and ideate early on and formulate them into a plan. **Product management** is a creative job first and foremost, and the most difficult part is knowing where to start. The beauty of starting with a product that already exists is that you have a starting point. Once you're able to get started with some of the lists and exercises we've already gone over in this chapter, you're ready to start building your product strategy.

This will serve as a high-level plan for how your product will be built over time and incorporates the many facets of expressing a product: the marketing, the distribution channels, and how it will be communicated to your customers. Keep in mind that many of the elements for your product strategy will spring from the first stage of the product development process: *ideation*. During this stage, many ideas will be explored and each of them will have a priority. These ideas will inform where you currently are with your product and where you want to go.

Reconciling all of this will be an act of putting all these pieces together in a rough draft, which will allow you to get to the heart of what building a product strategy is all about: deciding what your organization wants to achieve with your product and creating a product strategy that aligns with your overall company vision. If a product strategy is the "how" to the "why" of product vision, this means it will describe how the team will get to that ideal future state. The product strategy will break down the vision into actionable steps and priorities that will ultimately guide the roadmap and established product goals. It will include things like:

- How the product team wants to compete in the market and position the product
- Whether customers will be targeted and segmented around the product
- Key differentiators in the market and key features to build

Product goals

Next up, you establish goals for your product – you can't really have a strategy without establishing high-level goals. These goals will be a critical thinking task that will spring from your product vision, particularly now that this vision encapsulates the adoption of AI in whatever form makes sense for your product. These will be actionable checkpoints derived from the product strategy and they will provide specific metrics or outcomes that indicate success has been reached toward the product vision. Because they're goals, they will be short-to mid-term targets that will help indicate progress is happening toward that overarching product vision and strategy.

Goals exist to help you measure and prepare for important objectives that you want to reach with your product, for your customers, users, market, developers, and the rest of the employees in your organization. Consider these your strategic **objectives and key results** (**OKRs**) for your product. Without putting these big boulders in and building your product vision and goals from the place of what's most important, you'll drown in the weeds and you won't build strategically. You'll be left chasing opportunities unintentionally and it may tempt you to just follow what your competitors are doing or what your customers are asking for. Given the importance of the move to AI, take the time here to begin from a place of strength. If you find yourself weighing feedback from your competitive analysis, your customers, or your internal teams disproportionately when finding this vision and establishing your product goals, think critically about whether that's actually appropriate.

Feel free to also make this process collaborative and jot down all the goals that occur to all your various stakeholders so that you can prioritize them together. Again, you can always revisit these goals and scrutinize them later, but try and see what you can come up with together, refine that list, and try to stick to three or four higher-level product goals here so that you can begin to approach your product roadmap with confidence. All these goals will manifest in being broken down into smaller bits, and those bits may take weeks, months, or even years to fully realize. That's OK. Your product goals shouldn't be something that can happen in days and weeks. They should be encompassing and strategic enough that they are long-lasting because they're there to inform your overall product roadmap and to ensure everyone is working in service to a unifying goal.

When setting up goals, remember the following critical questions:

- Does the goal align with our product vision?
- Are there certain major milestones or benchmarks that we're hoping to achieve with AI being incorporated into these goals?
- Is this goal addressing the most critical customer pain points?
- How does this goal help move us closer to our long-term vision?
- Which metrics will we use to measure success?
- How will achieving this goal improve the overall customer experience?
- Is the goal realistic and achievable given the resources and time?

The product roadmap

Finally, we've made it to the product roadmap. This is the last step toward establishing our product strategy and it's the most detailed part because it's meant to capture initiatives that have come from all of the strategic goal-setting and vision exercises. This is where the magic of building and shipping products is. Taking strategic nebulous goals and turning them into a tangible reality, the roadmap is where we apply the conditions, limitations, and dependencies that we haven't yet bothered with. Roadmaps will change and they're not meant to be crystal balls, rendering your fate non-negotiable. They are meant to serve as guides for a product team's strategic directives. They are meant to be a bird's-eye, calendar-view translation of your product vision and goals.

Virtually everyone in your organization is going to see your product roadmap and it will serve as a blueprint for all teams to use to align when it comes to their own goals and planning sessions. It's also a way to communicate the results of all your strategy sessions and is meant to serve as a guardrail for all kinds of projects, initiatives, and implementation issues. Let's look at the key benefits:

- **Identifying changes in deadlines or scope**: If deadlines, issues, or bottlenecks arise, as they inevitably will, your internal teams have an internal source of truth to give them guidance on how much leeway they have, and it also gives them structure around when to escalate things because a roadmap is marked by months and quarters. This is an especially important note with AI products because, as you build, you will start to see certain experiments or projects taking longer than expected. Your regular product strategy sessions will make sure to absorb the shock of any major changes to deadlines or the scope.
- **Aligning different teams**: Establishing a product roadmap will also ensure you're aligning the efforts of your product design, development, UX, success, marketing, and sales with your strategic goals. Each major initiative you include in your roadmap will have its own documentation that comes with it in the form of user stories or product requirement docs, and all major initiatives will come with some kind of acceptance criteria. This makes it so that your roadmap is tied to the manifestation of specific outcomes. The roadmap isn't just a collection of features you plan to deliver by a certain date. It's a representation of the outcomes you expect to see from your development efforts. Multiple features may be in service to these outcomes.
- **Aligning the product with vision, goals, and expected outcomes**: As you start to establish a backlog of epics and tasks, you'll start to align them with the themes present in your product goals and initiatives as they get further grouped into sprints and releases. The product team, along with the AI PM, will own the product roadmap creation and maintenance and will continuously break down elements of the roadmap into individual tasks with their own acceptance criteria. Themes will start to emerge as you go. The roadmap will be a reflection of the initiatives and outcomes that are most compelling, whether it's because they have the biggest impact on customers or whether they offer the greatest ROI to the business.
- **Ease of working with stakeholders**: The last thing to keep in mind about establishing the roadmap is that, as an AI PM, you'll take whatever version of your product roadmap is final enough and you'll begin the work of evangelizing about that roadmap to your greater organization. Up until now, all of this has been a collaborative effort, but you can't get very far if your group of collaborators is too wide. Establishing a stakeholder team to work with that's representative enough of leadership, development, product, and go-to-market teams is a strategic task in and of itself. They will help you get to this point in the process. Now that you've gotten to this point, you're going to spread the word and get feedback from the rest of the organization because everyone will use it as a guideline. This offers you another level of scrutiny and allows for additional voices to raise concerns that might not have come up with the core stakeholder team.

To achieve this, you can maintain a recurring product strategy meeting with key stakeholders. We recommend making this a biweekly or monthly process if you're in a relatively new organization, and making it monthly or quarterly if you're at a more mature organization where things move more slowly. Make sure you're crafting your roadmap in a way where you're planning at least two quarters out so that you're strategic with your resources and time. Keep lines of communication open and allow space for new ideas and alterations to make their way into your product strategy and roadmap. This is a responsible way of not just capturing all the new energy that will come from building through your plans as they stand but is also a way to regularly welcome new perspectives. It's easy to get really great ideas when you're always brainstorming. When you're in the practice of just executing, brainstorming becomes a luxury that your time rarely affords, and things get stale. Resist the temptation to make these meetings ad hoc and only bring the party together when ideas start to stagnate. That's not a good practice because it puts a lot of pressure on your stakeholders to perform when the well has already run dry. Keeping this meeting scheduled regularly will make sure that your stakeholders expect it and get used to coming to it with their own ideas. Keep the meetings up, and even if you do run out of ideas or motivation, you can always shorten them when they do arise and give yourself more time back.

Now that you have a product strategy and roadmap you can be proud of (yes, this is the intended final state), we have great news for you: it's never going to be finished! This is just the first iteration of the current state of affairs. The goal here isn't to finish anything, it's to agree on the starting point. Over time, your strategy and roadmap will both evolve and you will revisit them regularly.

Remaining committed to keeping a standing meeting means that your stakeholders can take it upon themselves to keep a running list of topics and considerations to bring up in these meetings. It also keeps you honest by keeping you out of the very understandable and common pattern of staying in execution mode. When we're there for too long, our vision can get blurred by the inertia of building products. Maintaining the product strategy is your ultimate responsibility and it consists primarily of the strategic elements that happen when building your product vision, goals, and roadmap. Having regular product strategy meetings makes you a good steward of your strategy.

Red flags and green flags – what to look for and watch out for

Moving forward into the wide frontier of AI products includes some common pitfalls in AI transformation, as well as markers for success. We will call them red flags and green flags and they're signals you can pick up on as you go through this process of getting your product ready for AI adoption. Some of these will be concrete actions or results and some will be more emotional, but either way, you can use them as markers to know whether you're on the right path or whether you're hitting potential rough patches. In the following subsections, we'll address a few red and green flags. Let's get started.

Red flags

Red flags are behavioral patterns we can try and look out for that indicate there is an issue with a process we're setting forth. Because AI adoption is such a revolutionary undertaking for any company to take on, it's best to look out for some of these habitual patterns early so that you're starting on the right foot:

- **I don't know**: If you're regularly engaging with various teammates or employees on different teams at your company and they seem confused about why the company is adopting AI, what the value will be for their customers or readers, and what the product's high-level goals are, that's a red flag. The point of crafting a product strategy that aligns with the move to embrace AI is to make sure no one in your organization is confused about why this is happening. Make sure the entire company knows why you're adopting AI and which use cases you're trying to solve with it.
- **Shifting goals and strategy**: Building a product strategy or overarching goals is about setting an intention and sticking to it. The bigger a ship is, the slower it gets moving, so if you're finding the goals and strategy are often changing from the top down, it will make for unfocused, unaligned teams handling day-to-day tasks. Likewise, if daily tasks stray too far from the original goals and strategy, mechanisms have to be put in place so that teams are reinforcing higher-level goals in a more focused way.
- **Communication stalls**: If you're not meeting regularly or there's a lack of engagement in your recurring product meetings, there's likely some internal resistance you have to deal with. Do your best to communicate the exciting opportunity that embracing AI is, not just in terms of your product but the market as a whole. Rather than seeing AI transformation as a burden that's going to be a lot of work, you'll know you're on the right path when you can sense curiosity and openness from the stakeholders who hold leadership positions at the organization they serve. A lack of executive engagement is a critical red flag. AI isn't something leadership can outsource. They need to be involved and engaged at every stage of your AI adoption journey.
- **You're not seeing progress**: If you've started the process of working on creating your AI/ML pipeline and you don't see progress or at least some early positive signals from the model training you're starting, you might not have the right skill level to truly tackle the undertaking that is AI transformation. Making sure you have the right talent team, tech stack, and know-how is hard to do at the outset, and many companies struggle with establishing the right combination of factors as they get started. If you are struggling for talent, consider working with AI consultancies that can help you leverage from a diverse talent pool that's had exposure to use cases similar to yours.

Green flags

In this section, we wanted to include a few green flags that might appear as red flags at first glance. In part, we want to prepare you for some of the hurdles you'll encounter on your journey to help manage your own expectations, but we also want to convey that the following behavioral patterns are often a sign of things going right even if you might not experience them that way:

- **Your pilot project fails**: It was probably going to anyway. We've talked so much about managing expectations, and this is really what it's all about. Whether you're balancing your own reaction or the reactions of others in the face of lackluster performance or no performance at all, you're going to have to act as a source of faith for yourself and others. ML works and companies all over the world use it constructively and productively every day. Those with the patience and determination to see it through will be rewarded, but it can be hard to see progress, particularly at the beginning and when you're starting from the perspective of a traditional software product. Give it time and don't give up.
- **Too much feedback**: Whether from your customers, prospects, users, or internal teams, if you find yourself inundated with feedback about what should make it to the roadmap and what shouldn't, take it as a sign that you've succeeded in properly evangelizing your product vision, goals, and roadmap. When it comes to AI transformation, having more voices involved early on can be exhausting, but it's also an amazing indication that people are engaged. It's a good problem to have. Being the referee is easier than being the cheerleader who's trying to muster up excitement where there is none.
- **Reimagining data consumption**: If your internal teams are changing how they use and ingest data, that means the culture is changing toward becoming more data-driven, which will further support the success of your AI platform. If your organization realizes it has to massively overhaul how it annotates, stores, collects, and uses data, that means the AI transformation has gotten to a phase where it's affecting internal teams and the predominant climate.

As the AI PM, you'll want to be aware of the potential hurdles to look out for, as well as the indicators that your AI adoption is going well. Your organization will be looking to your leadership as this journey unfolds. Your credibility in this area will be cemented if you can be a safe haven for your colleagues when things go wrong and a source of encouragement when things seem to be going wrong but aren't.

Summary

So much of this chapter has been about how to anticipate and prepare for the jump to AI/ML. In *Part 2* of the book, we heavily discussed concepts related to the AI-native product: a product that's created with AI initially. Once you do make the jump to fully embrace AI in your own product, you can refer to *Part 2*, which is more focused on the aspects that come up when you're in the flow of building AI. In this chapter, we wanted to focus on the preparation stages for embracing AI/ML because of the gravity that comes with AI transformation.

Brainstorming ideas, vetting those ideas with practical considerations, getting your data right, evaluating the competitive landscape you'll be playing in, and bringing in your stakeholders to make a plan for how to build the transition are all part of AI readiness. All of the ideas expressed in this chapter are also easier said than done, and each section of this chapter will be a process in and of itself, but once you are able to get to the point where you've decided on a product strategy and you're executing that strategy with a product roadmap that your entire organization can get behind, you're on your way to building with AI.

It's a privilege and a massive responsibility to be a PM, particularly one that oversees the AI transformation of their product. It's an exciting time in your product management career and one that you'll likely look back on fondly. We're at a unique precipice in history and not many PMs can say they were able to be of service in such a big step for their organizations and their respective products. As you hit difficult moments along this journey, remind yourself that this is a difficult transformation to manage and evangelize for others in your organization too.

Embracing AI and all the challenges that come with it isn't for the faint of heart, and if you find yourself getting frustrated, it's actually a great sign. It means you care enough about your product and your organization's success with AI to elicit an emotional reaction. Give yourself and your co-creators grace and be proud of the work you're all doing collectively to embrace the wave of AI.

In the next chapter, we will explore the role of design in more depth. If this chapter sets the stage for AI transformation, the next chapter shows us how products can evolve to embrace AI from a practical design perspective.

Additional resources

- Udacity's *Intro to Artificial Intelligence course* (`https://www.udacity.com/course/intro-to-artificial-intelligence--cs271`) and *Artificial Intelligence* Nanodegree Program (`https://www.udacity.com/course/ai-artificial-intelligence-nanodegree--nd898`)
- Stanford University's online lectures, *Artificial Intelligence: Principles and Techniques*: `https://stanford-cs221.github.io/spring2022/`
- edX's online AI course, offered through Columbia University: `https://www.edx.org/course/artificial-intelligence-ai`
- Microsoft's open source Cognitive Toolkit (CNTK) to help developers master deep learning algorithms: `https://learn.microsoft.com/en-us/cognitive-toolkit/`
- Google's open source TensorFlow software library for machine intelligence: `https://www.tensorflow.org/`
- The **Association for the Advancement of Artificial Intelligence** (**AAAI**)'s **Resources** page: `https://www.aaai.org/Resources/resources.php`
- *A visual introduction to machine learning*: `http://www.r2d3.us/visual-intro-to-machine-learning-part-1/`
- *Machine learning with PyTorch and Scikit-Learn: Develop machine learning and deep learning models with Python*: `https://www.packtpub.com/en-us/product/machine-learning-with-pytorch-and-scikit-learn-9781801819312`

- *State of AI, Chapter 7*: `https://www.stateofai2019.com/chapter-7-europes-ai-startups/`
- *8% of EU enterprises used AI technologies in 2023*: `https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20240529-2`
- *100+ updates that make Shopify's foundations even stronger*: `https://www.shopify.com/news/edition-winter-24/`

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

15

The Role of AI Product Design

The decision to integrate AI into traditional software products is two-pronged today. It gives you a competitive edge in your market and it also helps you better understand your users on a deeper level. We've discussed the concept of AI integration at length in this book so far, and the influence this has on **product design** can't be understated. Incorporating AI isn't just for those who want to "get ahead"; it's also for those who want to remain relevant in their playing field. Incorporating AI also allows companies the opportunity to better understand their users and customers and, in turn, make their products even better.

Virtually every company out there today will be incorporating AI into their operations, internal workflows, product experience, and strategic analysis. This chapter will focus on the design process when building AI features and capabilities in your product experience. Before we can manage the day-to-day operations of a traditional software product, we have to discuss how AI changes a number of areas and sets the stage for a newly evolved, AI-powered counterpart to emerge. This transformation process isn't one that should be taken lightly, and this chapter will go over the many aspects of AI transformation that deal with how we ideate, develop, and communicate the value of this emergent product 2.0. By the end of this chapter, you'll understand how you can manage the many facets of product design, communication, and development when working on an AI product.

We will cover the following key topics:

- The evolution of product design
- What makes the evolved AI product special
- Choosing your words carefully
- Building with trust and security

The evolution of product design

Product design is not a limited, siloed discipline of its very own. Sure, there are important considerations you're making when deciding if a particular button should be blue or red, or when you're analyzing how many keystrokes it takes to get from point A to point B in your product. When many people think of product design, this is what they think of. But that's just the tip of the iceberg.

Product design is more encompassing than the visual layout of your UI. It's a crucial part of product management that affects everything from ideation to deployment. Design decisions can happen at all levels of the AI product development lifecycle. Recall the stages of the AI product development lifecycle from *Chapter 6*, as shown in the following diagram:

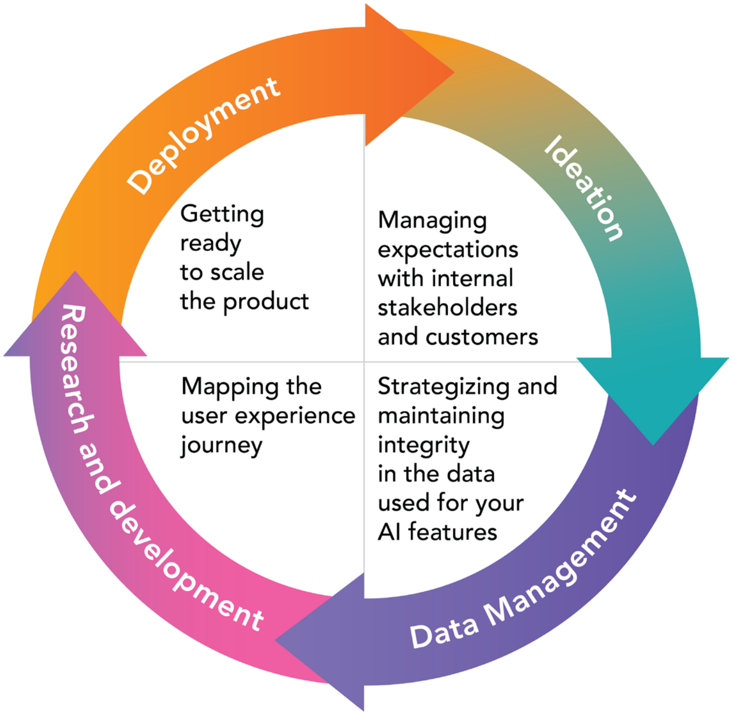

Figure 15.1: AI product development lifecycle

We will be referring back to each of these stages as we make our way through the following sections. Incorporating AI into your product isn't just about a straightforward enhancement; it represents an *evolution*, that is, a fundamental shift and re-evaluation of the product's design, user experience, and functionality. Returning to the drawing board as a logical first step when contemplating what AI integration looks like is about ensuring that product teams conceptualize the overall product value proposition and user experience with AI, rather than adding capabilities casually as an add-on. The final output shouldn't be an AI feature that does something marginally better. It should be a re-invigorated, reconceptualized product experience that's substantially better because of AI.

Here, we will be focusing on the importance of design at every stage of the product development lifecycle. Let's start with how we can manage expectations.

Ideation: Managing expectations

In order to bring a product to market that's significantly better than its first-gen counterpart, we have to understand the expectations and experiences of our users' needs. These needs are likely to change over time as the competitive and technological landscapes change as well. Theoretically, a product team is already aware of their users' issues, concerns, pain points, and demands. After all, maintaining a feedback loop with your customers is already a foundational part of a product role at most companies.

If you've been on a product team managing a traditional software product for some time, you've likely already got a baseline for how your product is currently serving your customers. You have a backlog of potential features you're already planning to develop. Even traditional software products have predefined algorithms that help to solve specific problems based on user inputs. The integration of AI doesn't necessarily change any of this, but what it does do is force us to re-envision what a product experience with AI might look like for potential customers. The existing version of your product is effective to some level, but it's likely that it falls short in other aspects. AI introduces the ability to fill these cracks in our product offering, turning shortcomings into opportunities for a renewal in product design.

Beyond the level of existing product functionality, new expectations and assumptions will come from incorporating AI. Today, AI is a polarizing topic since people have varying viewpoints. Some are excited about what AI might mean for them, while others are uncertain about what an AI-heavy future might look like. Some are disturbed that AI is learning more about our behaviors than we realize, while others find a world where AI might anticipate our next need a welcome relief. For most, AI can elicit conflicting feelings across different times and contexts. Either way, incorporating ML and DL models will redefine what your users and customers come to expect from your product.

Taking a step back after the strategic decision to incorporate AI and go back to your early design principles allows product teams, developers, data teams, leaders, and UX designers the chance to reassess the core problems your product is offering a solution for. For example, if your product is a customer relationship management platform, perhaps it's gotten really good at managing customer data and logging interactions. But with AI, your product could learn from those interactions. It could offer your customers a sense of partnership, helping them with predicting customer behaviors and future interactions, automating outgoing messages, or offering advanced analytics and insights to your users.

Now your product isn't just fixing one problem it already knew about; it's helping your customers and users anticipate future problems they weren't yet aware of. This brings the product into a new echelon. If your customers accept your pricing model and value proposition in comparison with the other players in your market under the conditions before, what does that mean for future iterations of your product with new AI capabilities? Then, you get into a different kind of conversation. Does this new iteration of your product give overblown expectations to your customers about what they can reasonably expect from your product?

This is true of your internal stakeholders as well. Part of your work as a PM, particularly early on in the AI ideation phase, is to manage the expectations of the key stakeholders that also come with their own assumptions and notions about AI. Balancing expectations on the customer side and internally isn't easy. It's one of the reasons it's recommended to include stakeholders in product design discussions and brainstorms as early as you can. Doing so will mean you get ahead of their input, perspective, and concerns, but it also means you can soften the blow of a paradigm-shifting technological adoption.

Your development teams will also be impacted by the integration of an AI system as it relates to scalability because of the impact it will have on development in general. In traditional software development, we're often working in some kind of Agile framework that follows a linear path of gathering requirements, designing, developing, and testing. But this process will be stressed with the implementation of AI because of the iterative nature of AI. Though an organization may already be agile, the development and management of AI systems in your existing workflows will still have an impact. This is because AI model development involves more iterative cycles of training, testing, and tuning, which will lead to a slower time to delivery than what an agile engineering organization might already be used to.

The team will also need to start working with new kinds of infrastructure and tooling, like data pipelines, model training environments, and ML frameworks, once AI is involved. Your existing development process and culture will be impacted by experimenting with various models, training methods, validation methods, and model tuning as you are operating within the confines of your established development cycle and legacy systems. Adopting methodologies to catch the continuous evolution of your data science and machine learning functions and pipelines will ensure you're setting your development teams up for success as they stretch to accommodate the new paradigm of AI.

To successfully manage the design expectations of all your internal and external stakeholders, you should be able to answer the following questions at the end of your ideation phase:

- Can your stakeholders clearly define the core challenges in your product that AI will address?
- Can they articulate how AI will improve the product for users? Which specific pain points will it alleviate?
- Can stakeholders articulate how the new AI features and capabilities will differentiate your product from competitors in a new landscape?
- Do all relevant business and technical stakeholders understand the measurable goals and metrics tracked for the new AI features and capabilities?
- Are all stakeholders aware of how the new AI features and capabilities relate to product strategy, vision, and goals? To company goals?
- How will the product interface evolve with AI and how will users interact with AI-driven features?
- How will the user journey and customer behavior change with integrated AI features and capabilities? How will this be communicated to customers and users?
- How will the integration of new AI features and capabilities impact existing business processes and operations?
- What are the expectations around re-training or upskilling staff to better adjust to this impact?

Data management: Strategizing and integrity

Because AI products are so reliant on data, there will have to be significant discussions about what that means for your data strategy overall. How will data be entered, transformed, processed, stored, and trained on? What system will you put in place to ensure the quality remains high, that is, that it's safely stored and retrieved? **Data architecture** and **data management** are dedicated practices on their own, and many companies can find themselves overwhelmed by the data demands of building AI features if they're not prepared for the cost and competency of understanding these considerations before they've deployed any AI features.

Part of going back to the drawing board means necessary conversations about how the data will be captured, categorized, and processed to make sure it actually makes your AI capabilities useful. Testing your AI features to make sure they work as intended will often come down to how reliable and scalable your data is. We will discuss AI bias and ethics later on in this chapter from a design perspective, but this is a consideration that's worth noting in this first section because you'll likely want to have some alignment on how you're going to address data security early on in your AI ideation process.

While you might not necessarily want to have all these decisions made at the ideation stage, you should be setting aside dedicated time for these discussions to happen early on because the determinations you make will have a downstream impact on how your product is built and developed. It will certainly have an impact on how it will amass storage and compute costs. Finding a way to be agile with your ideations and start vetting requirements, costs, and scope for each, even if they're rough estimates, will help you plan strategically from the beginning.

To successfully manage data strategy and integrity, you should be able to answer the following questions at the end of your data management phase:

- How will data to support your AI features and capabilities be centralized and structured for AI/ML?
- Have you established processes to execute data preparation standards to handle your proposed AI features and capabilities?
- How will data pipelines be monitored to remain reliable to process new volumes of data?
- How will quality issues in data or data pipelines be flagged and addressed?
- What kind of robust privacy measures will you plan to carry out in order to make sure your customer and user data is stored appropriately?
- Are you planning to make transparency and data privacy a central fixture of your product offering in the first place?
- Which data governance and ownership standards will you implement?
- How will you address issues of data bias? Will this be something you test for regularly?
- How will you make sure training data for your models is accurate and up to date?
- How will your organization comply with national and global regulations like GDPR, CCPA, and others?

R&D: Mapping the user experience journey

If you've been on a product team for some time, you will already have your user experience journey mapped to some degree. The level of formality of this depends on your company and product management style, but most PMs will have a deep understanding of where they might lose users along the journey. In many cases, these insights will guide your backlog prioritization strategy and will inform how you plan ahead in your roadmap. Some PMs find it tempting to keep the existing UI and UX flow of their product experience intact and add AI incrementally, but often, integrating AI demands an overhaul of these existing user journeys.

The reason for that is an AI-enhanced product is introducing new interaction paradigms. Rather than your product being fixed, with specific inputs that offer deterministic interactions, AI products are dynamic and are often learning from contextual interactions. Incorporating recommendations, analytics, or NLP in the form of chatbots will mean your users will interact with your product differently. You'll also likely want to build these features out in a way that feels intuitive for your customers and users.

These new interaction paradigms allow a product team to base their research and development stage of the AI development lifecycle on identifying pain points and opportunities along the user journey. Focusing on the journey means you can think about how AI could have an impact on the design of your product in a way that limits the pain points, frustrations, and bottlenecks your users face. This means that R&D activities will be based on how new AI features could align with user needs in a way that enhances user satisfaction and, ultimately, engagement.

Perhaps AI also offers your users a substantial difference in the way your customers interact with your product. Maybe, instead of typing and clicking on things, your product requires voice or conversation. Perhaps you're going to need new visual elements to make interactive dashboards where users can simulate various scenarios. Perhaps, instead of categorizing inputs themselves, your users will instead have to gauge the efficacy of an AI system and serve as reviewers.

The use cases of AI are still quite limited. You're likely helping to automate, predict, or recommend something in your product experience. Or perhaps your AI features are more generative in nature. Either way, most of these use cases will require your users to engage in a vastly different way from how they were interacting with your product prior to its AI evolution. If you can create a map of your UI/UX journey with some known AI features you'd like to test early, you'll be better prepared for the significant impact AI adoption will have on your customers as well as the investment AI implementation will require of your design, product, and development teams.

Some of the key R&D design considerations when mapping your user journey are:

- **User centricity**: Define and document how certain model choices or solutions are impacting users directly to ensure the focus remains on user experience. Can there be ways your new AI features and capabilities can adapt to user preferences or create more accessibility among all users?
- **Touchpoints**: Identify specific areas where AI/ML models can have the greatest impact when interacting with users. Make sure those touchpoints integrate with the existing user flows that have already been established to minimize the user's learning curve.

- **Explainability:** Define how you will manage expectations with users about the new features you're targeting and what level of explainability you want to pass on to customers and end users to enhance transparency, trust, and adoption. These will give you the blueprints you need to eventually create onboarding tutorials and feature guides.
- **Bias:** Establish ethical guidelines and mechanisms to ensure fairness and inclusion during the R&D phase to further enhance trust in AI outcomes and outputs.
- **User feedback:** Build processes that allow you to gather feedback from real users from beta testers and other select groups of users to make sure you're fine-tuning models appropriately at the level a typical user would expect.

Deployment: Are you ready to scale?

Often, when we talk about scale, it's about whether a product can technically scale up to support a 10x or 100x increase in customer acquisition. Technical debt and technological limitations are one thing, but costs are another. Because AI can be computationally heavy and data-heavy, some of your earliest considerations for your product might be: is it worth the cost of offering this piece of functionality to users? Is the feature strategic enough that it's worth the resources, financial investment, and technical debt to support? Remember that technical debt is a form of investment that will eventually need to be paid off. Acquiring technical debt might allow you to release new features faster in the short term, but if left unaddressed, it can accrue interest and cause bigger problems in your AI system later.

Once you're in the deployment phase of the AI product development lifecycle, you're integrating models and data pipelines into legacy systems, finalizing your infrastructure for go-live, working on a rollout strategy for the new AI features and capabilities, and setting up mechanisms for the continuous maintenance and monitoring of your new AI system. Theoretically, if all goes well, you're supporting the existing volume of users as well as the new users you will acquire in response to your AI investment.

The way you grow and scale as a result of AI adoption could result in new volumes of data you weren't gathering before, which could have implications on the costs of maintaining the AI system you've built to support new AI features and capabilities. Likewise, your thresholds for performance will impact price and scalability as well. Perhaps certain accuracy thresholds are just too expensive to justify than others. Your data science teams will be coming to you with a range of options. Likely, it will be the models themselves they're comparing performance against, but it will also be performance thresholds that need to be understood and discussed as well. Going back to the drawing board with regard to scale is relevant here because there might be infrastructural decisions that need to be made about how you'll support your code base and data accessibility for AI integration.

Some of the key questions you can think about when you're in the deployment phase that relate to scalability include:

- Could cloud-based platforms and tools offer you a way to beat the new margins AI integration will demand?
- Are there ways you can limit computational demand or use distributed computing frameworks to help serve your ability to scale?
- Are there hard limits on the number of customers you'd be able to support?

- What refinements can you make to your data infrastructure to make the retrieval and refreshing of data more cost effective?
- Could you add more servers or upgrade existing servers to minimize latency and meet performance thresholds?
- How will model versioning and continuous deployment be handled in a way that new models can be deployed seamlessly as the application or feature grows?
- Are there strategies like load balancing, distributed processing, or optimization techniques that can be used to reduce strain on AI/ML models in deployment as user adoption grows?

Now that we've covered the various stages of the AI product development cycle and how they apply to traditional products evolving with AI functionality, let's explore potential routes AI PMs can take to make the most of AI.

Expansion: What makes the evolved AI product special?

If we're not strategic about the impact and value AI capabilities will bring to a product experience, we're investing in an infrastructural and resource-intensive program that won't bring the reward we envision at the beginning. Evolving an existing traditional software product into an AI product isn't a paradigm shift only in terms of what experience it will offer your customers and users; it also impacts how your company effectively operates. It's an internal and external paradigm shift that's felt by all.

That's why being intentional with how you're going to approach AI transformation is important from the start. You want to end up with a product that's significantly improved with AI because the alternative is costly in many ways. In this section, we'll talk about the aspects of AI integration that can bring value and the various ways that value is manifested.

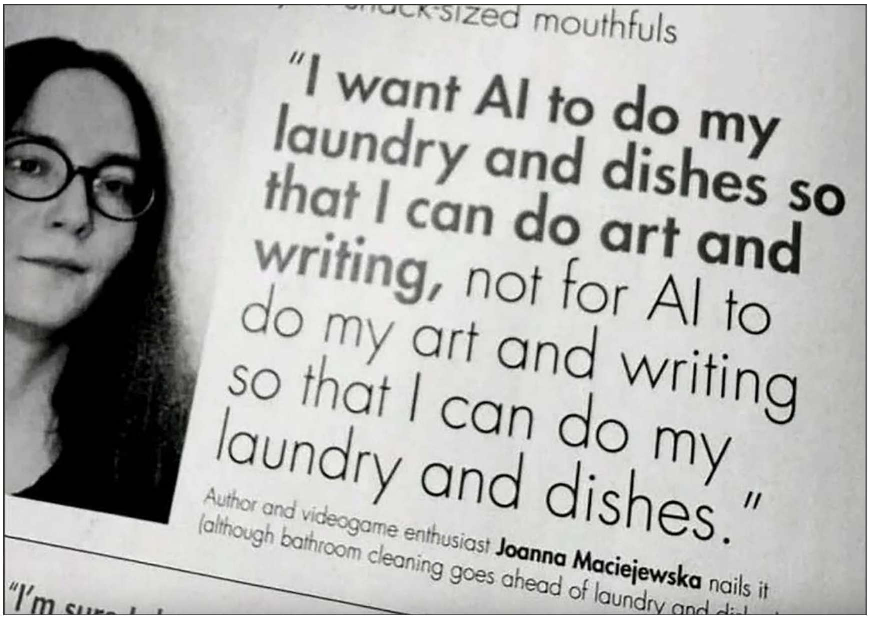

Figure 15.2: What is AI good for?

When we think of AI adoption, most people would appreciate it if AI took on aspects of our everyday lives that we actually *don't* want to do. The above quote encompasses this sentiment most elegantly, and it's in the spirit of Joanna Maciejewska's quote that we consider how to best think about what value AI may bring to our product. Do the proposed AI features we want to incorporate help people feel like they're being helped and supported with the drudgery of their days or workflows? Because if not, that feature likely isn't bringing in that much value. In the following sections, let's take a look at some of the most impactful ways AI can be integrated into a traditional software product and how that might impact your product design.

Decisions and insights

Because you'll likely need a lot of data to run your AI systems, that will mean that you'll be gathering enough data to offer significant insights that will impact the decision-making being done on the customer end. If you have enough data to run an AI system, you have enough data to find patterns in it and help your users make decisions in a way that's more informed. Analyzing habits, predicting future states, offering recommendations, or forecasting future costs and trends is a way to use insights collected from your users' data that are meaningful to them. You can do this at the local level, by analyzing the data at each user level, or you can append this data to external data sources.

From a product design perspective, this gives you many options to choose from when it comes to reimagining your product experience. You can build interfaces and experiences around your product that demonstrate to your users that you are learning from their behaviors and preferences. Not only that, but you're using that information to in turn help them with their own decision-making through dynamic interfaces, visualizations, and dashboards. For example, Salesforce integrated AI into its existing **customer relationship management** (**CRM**) platform through Einstein AI, which helps users working in sales to make decisions more effectively. Einstein AI allows users to receive predictive insights about sales opportunities, customer behavior, and marketing campaigns so they can prioritize leads and forecast deal closures more effectively.

By giving users and customers the control and flexibility to understand more about themselves, their environment, or their peer group, you're showing them that they can trust in the product experience you're crafting for them. We don't look at trends, insights, and forecasts in a void. We care about these things because more knowledge empowers us to make better decisions in the moment. If your product can show even one user that they're better prepared in an uncertain world with your product in their arsenal of tools, you can win their loyalty. If you can do that in a way that presents decision-making and insights in a filtered, digestible format, you can win their admiration.

The key considerations for ensuring that you have enough insights for effective decision-making for product design are:

- **User personas**: Build personas around insights that assist decision-making for different sets of users. Not all users will appreciate the same insights. Perhaps some users want more in-depth data while others want more high-level trends. Find what is most valuable to your users and tailor insights that are contextually relevant and actionable.

- **Collaborative design**: Engage your internal stakeholders through workshops devoted to developing insights and analytics features to make sure you have holistic input on which insights would be most meaningful and how they should be presented to maximize user satisfaction and expectations.
- **Testing**: Build mockups and prototypes of visualization and insight features to test with real users to make sure you're getting feedback early. This will allow you to iterate more effectively and ensure final designs are relevant and user-friendly before a massive rollout.
- **Continuous learning**: As you add more complexity with regard to data, insights, and functionality that informs decision-making, build training materials and onboarding resources that will help users understand how to interpret the insights and visualizations they see.

Automation and adaptability

Automation is one of the easier AI features there are to work with. Taking considerable time away from an arduous task that someone no longer needs to perform themselves is a surefire way to demonstrate AI's efficacy and value. In fact, automation is one of the aspects of AI that people think of first when they think of AI value because it can reduce so much cognitive load on users. But it's often an over-amplified expectation that rarely plays out seamlessly in reality. So, if your product can offer a seamless automation experience, you're already ahead of the curve.

Automating a task that reduces time or cost delivers quite a lot of value already because it eases the burden on users. From a product design perspective, this means prioritizing features that have a direct impact on the usability and efficiency of your product. Automating in a way that enhances someone's experience of your product and offers a superior user experience surprises those with the highest AI expectations and further contributes to user satisfaction and retention rates. Adapting your product experience in a way that learns from the snags that come up during automation will keep your customers engaged and anxious about your next release.

For example, customer support software company Zendesk incorporated AI by adding a chatbot functionality into their platform called Answer Bot, which helped automate responses to common customer questions by learning from previous interactions and adaptations over time. Because it was able to learn from previous examples, it was able to evolve answers when new discoveries or considerations were uncovered. The automation reduced the need for routine and manual tasks and helped live agents focus on more pressing, complex customer issues.

To ensure that you're automating the right way, here are some questions you can ask of product design:

- Which repetitive or time-consuming tasks that users frequently perform can be automated?
- Which automations will not only save time but also improve the user journey?
- What are all the data types and sources necessary for an automation feature?
- How will automation features be designed intuitively and clearly for the user to understand and engage with?
- How will users report issues or suggest improvements for automation features?

- Are there potential avenues you can take to make automation adaptable to user needs and preferences over time?
- How will you gauge the success of your automation features and adoption?

Personalization and learning

We crave being known and understood. If your product's AI features and capabilities can demonstrate that it's learned from a user's behaviors, it will likely develop a stickiness with that user. Traditional software products struggle with achieving this level of personalization because they aren't built for that. AI can not only offer us a chance to know our users well enough but also convey that we have the ability to grow and learn with them as well. As the AI space becomes more and more saturated, the companies that can deliver a product experience that can nimbly adapt to users' preferences and behaviors will be the ones to benefit the most from their AI investment.

As user satisfaction and engagement grow over time, it will lead to higher rates of customer loyalty, lower customer acquisition costs, and higher conversion rates overall. For example, online learning platform company Coursera was able to enhance its user experience by introducing personalized learning paths so that they could tailor their course recommendations based on their learning patterns, interests, and historical data from previous courses. As users engaged with content on Coursera, these personalization features were able to better refine suggestions and help users discover learning styles and skills that aligned with their career goals, keeping them motivated to keep learning.

This relates to product design because it emphasizes the importance of implementing data collection mechanisms that allow you to gather the kind of behavioral insights that will enable personalization downstream. Reimagining your existing product in a way that emphasizes the development of features that respond to user needs, preferences, and goals will help create deeper connections between your product and users. Here is a list of questions you can ask to make sure you're making the most of your personalization options as you reimagine your product design:

- Which specific user behaviors can we track to help offer recommendations and personalized insights to users in their workflows?
- What kind of continuous feedback loops can we implement to help users understand and share their preferences with us?
- Where along the current user journey can personalization be introduced to significantly enhance user experience?
- Which user segments can be defined based on a set of user behaviors and preferences and how can our offerings be tailored for each segment?
- Are there current barriers or obstacles that would make personalization challenging and how are we trying to address those challenges?
- What's the potential ROI for developing personalization features?
- How can increased personalization be correlated to customer loyalty, user retention, and user engagement rates?

Choosing your words carefully

More and more companies will start to get their message out to their markets as they make the steady and sure transformation into AI. Getting messaging right with regard to your AI-evolved product can't be overstated because it's something companies consistently struggle with. As companies pivot to include AI in their product offerings, they will need to be thoughtful about how AI capabilities and features are communicated because misleading messaging can lead to disillusionment among existing users. AI PMs will need to ensure that the language used is reflective of the evolved product's true capabilities and manage expectations accordingly.

As a PM, I'm always interested in how companies articulate their AI offerings and value because I rarely see them getting it right. Most of the time, AI features are communicated as a silver bullet, priming their customers for an experience they're rarely going to experience in practice. Incorporating clear and precise language about AI features and capabilities early in the design process is essential because it helps bridge the gap between what users expect and what the product actually delivers. Integrating these considerations into the design phase means that PMs and designers can create a new product experience that resonates better with users.

For example, IBM Watson Health was heavily marketed as a leading-edge, revolutionary tool for diagnosing diseases and helping medical professionals by suggesting treatment plans for patients. The messaging was misleading and positioned the product as something capable of more than it was. According to Robert Wachter, chair of the department of medicine at the University of California, the company "came in with marketing first, product second, and got everybody excited". Hospitals found that the system struggled with real-world medical data and offered irrelevant or incorrect recommendations.

Being clear on how you want AI capabilities to come across to your customers early on is also a way to align with user perspectives and reimagine your new product experience in a user-centered way. Incorporating ideas of product language fit around AI features early in the design process is a way of ensuring you end up with a final product that won't surprise anyone and is built for them. Whether you sit in marketing or in the user's seat, you'll want the delivery and expectation around your AI-powered product to ring true and send the right message.

Product language fit

Creating strong product language fit means being thoughtful with your communication style and using precise word choices to convey the appropriate tone that will resonate best with users and also manage their expectations. Some aspects of product language fit and the way you talk about your product's AI features and capabilities are:

- **Communication style**: Keeping your messaging clear and simple goes a long way. You want your language to be accessible and memorable. Avoiding unnecessary jargon or overly technical language means that your general audience will find your messaging digestible and it will help you engage with a broader audience as you re-release your AI bolstered product. Keeping the product messaging on user needs and experience, rather than on the AI technology itself, is a way to emphasize and contextualize how AI features will be useful or relevant for users' real-world experiences, which will, in turn, make your product more relatable and valuable to your audience.

- **Word choice:** Look for wording that will convey the truth about your AI features and capabilities without exaggerating. For example, rather than using terms like "magical" or "revolutionary," you can reach for more grounded phrases like "powerful" or "advanced" to temper expectations. You'll also want to avoid hyperbolic words that convey extremes. For instance, limit how often you're reaching for definitive words like "always" or "never" when describing AI features and capabilities so that you're not saddling potential or existing users with unfair perceptions of what your evolved AI product can offer them.
- **Tone:** While many may be excited about the prospect of incorporating AI into their products, that excitement can prime users for future disappointment. You'll want your messaging to convey confidence and optimism for your product's next evolution, but in a balanced, informative way. Acknowledging the strengths AI can bring to your product experience, as well as the limitations of the new features and capabilities, is a way to remain transparent and relatable.

Product communications aren't just limited to marketing. Everything from your product documentation and thought pieces to social media comments, blog posts, marketing materials, and change logs will be impacted by decisions you make early on about what AI integration will mean to you as a product team and to your users as those experiencing your product. Building consistency with all these communication channels early on will foster trust with your user base.

Disparities between channels or using certain words over others in different contexts will only create confusion and distrust. It may even affect the perceived reliability, efficacy, and competency of your product. When products are positioned and communicated inconsistently, it sends the message that they were built inconsistently too. Building an AI program, supporting AI features and capabilities, and retraining and preparing your workforce for the shift it will bring are too big of an undertaking to leave communication up to chance.

You will not need to have all your marketing collateral ready early on in the product design phase. But you will need to have a devoted meeting or brainstorming session about it as soon as you can. It should be a foundational part of how you're ideating the value of the AI features you're proposing to add to your existing product. Even if product communications are coming from different sources internally, there should be a universally applied style guide that the product communications conform around. Launching a new AI-powered version of an existing product is big news for most companies at this point. This means it's also an opportunity to communicate in a way that's consistent with your brand identity.

As you continue down the product development lifecycle, you'll likely need to have a strategy on how you're going to talk about the ethical aspects of your AI features that deal with data privacy, governance, model bias, and training data bias. Establishing a relationship with trust early on will make you more credible if and when you are ready to divulge more details about your AI ethical philosophies. Even if you do plan to hold your AI sacred sauce close to your chest, users still need to know what the AI features and capabilities in your product can and can't do. They'll also need some information about how incorporated AI is processing and using their data. Building trust and setting realistic, achievable expectations is the path to launching a novel product in a way that's best set up for success.

Accessibility and inclusivity

Language plays a role in how features and products are understood and utilized. You're launching a new product once AI transformation has taken place and, once you do, you'll want your loyal base to come with you on that journey. Abandoning them with a product that doesn't look or sound anything like the existing one they're already used to will introduce a lot of friction to your product rollout and your bottom line. As you're designing your product features early on, you'll want to make sure the new AI features are designed and communicated in a way that maximizes a diverse user and customer base.

What this means is you'll likely have to incorporate ways of making your product accessible to different ages, backgrounds, technical proficiency levels, language skills, and even accessibility needs. Using clear and simple language that's understood by all is often the best way to make sure your product's value isn't misconstrued. For example, Microsoft launched its Seeing AI app to help visually impaired users understand the environment around them through image recognition. During the initial rollout, Microsoft hadn't involved the visually impaired community in the design and testing phase, which meant that the features they built didn't meet the needs of their users. As a result, the app struggled to recognize complex scenes or differentiate between similar colors. The UI was also not intuitive and needed simpler, clearer instructions.

Chances are there will be a few different profiles of various users that you're going to build product onboarding, materials, and resources around as you get closer to launch. You may not have these elements ready during the design phase but you should be thinking about them. Yes, you will be designing as inclusively as you can, but no matter how well prepared you think you are, there will always be edge cases. Even without edge cases, relaunching a product will always come with some sort of learning curve for even the sharpest of users. Building an onboarding library and knowledge base, tutorials, or modules if they're required is a way to make sure your existing and new customers are well supported to continue the journey with you. Here are some questions to help you with accessibility and inclusivity as you're evolving your product with AI:

- Are you demonstrating the value of your product in a way that aligns with what your users expect from your product already?
- How are the needs of various user profiles in your user base considered in your AI product re-design?
- Does the delivery of your AI feature messaging relate and contextualize well with the existing knowledge your existing users already have of your product?
- How will you gather ongoing feedback to improve accessibility and inclusivity in your AI-evolved product?

Building with trust and security

When evolving traditional software products into AI products, bias, explainability, accountability, and security considerations must be integrated into your design process so that you can avoid concerns in these areas after you re-launch. Bias can lead to the unfair treatment of users and unintended consequences that can hurt your brand and the product's reputation. Explainability in AI product design will help create more trust with your current and potential users.

Reinforcing accountability when relying on AI-driven decisions gives your users confidence when interacting with your product. Bolstering your product with security protocols and ethical design practices will also help mitigate the risks associated with the misuse of the data that will be used in your AI systems. Let's discuss these concepts in greater detail and the role they play in product design.

Bias

The perpetuation of bias with AI is a hot topic and it represents one of the areas we will need to focus on when addressing ethics and transparency with regard to AI. Work with your stakeholders early to understand how you're going to address bias, mitigate its impact, and put guardrails in place to ensure it doesn't veer too far. Whether or not maintaining trust and offering equitable product outcomes is baked into your company and product values, it will be an issue raised with your user base if the bias is strong enough. If left unattended, this will have an impact on your brand identity as well.

You can minimize the risks and impacts of this by talking with your data science teams about how you will manage bias and drift in your data samples. Beyond that, you can also conduct regular bias audits, train your models on diverse training data samples, and use fairness-aware algorithms. Whether or not you choose to incorporate this language into your formal product communications, these practices will still help you maintain control over your product performance and risk exposure. This is particularly true if your product has the ability to discriminate based on race, gender, or other protected classes.

Here are some key design-related considerations to ensure your product doesn't inadvertently have any biases:

- Include diverse demographic groups (age, gender, cultural, and socio-economic) to identify varying needs and perspectives in your user research (interviews, surveys, focus groups, etc.) when you can.
- Create prototypes and MVPs of AI capabilities early in the design process to conduct usability testing with diverse user groups.
- Incorporate bias mitigation strategies into design specifications/requirements for things like diverse training datasets, algorithms for countering bias, fairness constraints, and design patterns that promote inclusivity, like language options and accessibility features.
- Build feedback mechanisms that allow users to report bias or issues with the product directly within the product experience so they have a chance to tell you before they tell public channels.
- Keep cross-functional teams involved in product design diverse and make sure diverse viewpoints are welcome.

Accountability and explainability

Customers and users will want to know what they can reasonably expect out of your product, and depending on how AI-fluent they are, they'll also want to understand how your product arrives at certain decisions or determinations in simple, plain language. Whether your users are consumers or business users, you'll likely still have some explaining to do as a product team. Most users need to know that there is a clear line of accountability for the AI feature's actions. Proving to users not just that you can be trusted to build responsibly with AI but also that you can communicate how and why it is responsible will resonate with people.

To ensure explainability of your product, here are some aspects you must account for in your product design:

- Perpetuate design transparency by documenting and sharing how AI algorithms are being used through user guides, FAQs, and glossaries at varying levels of AI fluency so that all users can understand how the new, evolved product functions.
- Incorporate explanations or visualizations of the decision-making process that's happening with the new AI features and capabilities for users to see.
- Incorporate real-time explanations within the product's UI for users to interact with AI features more deeply where you can to make explainability a part of your product experience.
- Gather feedback regularly on how well users understand the explanations of AI features and capabilities through user testing sessions and make sure the design evolves based on user needs and comprehension.

Security

Because of the level of data involved, most AI systems pose a security threat. Securing these systems and the swaths of data that are used to power them is a considerable undertaking. Building safeguards that address both cybersecurity issues and AI-specific vulnerability testing and penetration testing is a good way to hedge against potential threats or sources of degradation for the AI systems you will build to power your evolved AI product.

You'll be working with a number of technical counterparts to create architecture that will power your product's new features and capabilities. Models themselves can be targets for manipulation and external attacks, so you'll want to design the system you'll be creating to detect and handle unexpected or malicious inputs that could skew your model's performance.

On the data side, you'll want to encrypt data both when it's at rest and when it's in transit between APIs. In some cases, attackers could try to poison your data by feeding your models data that can produce harmful or inaccurate model outputs. Depending on the features you'll be building, in many cases, they will be processing sensitive personal or business information, so preventing unauthorized access is important.

To ensure the security of your product, here are some aspects you can account for in your product design:

- Incorporate **role-based access control** (**RBAC**) to limit those who have access to various levels of data, AI/ML models, and outputs.
- Limit access and masking/anonymizing data when you're training your models.
- Embed security best practices into the development cycle with regular code reviews and checks for vulnerabilities both for the AI models that will power your AI features and the surrounding infrastructure.
- Build encryption into your AI systems for both data in transit and at rest.
- Securely deploy models using containers and isolated environments.
- Develop incident response plans for AI vulnerabilities.

- Help identify and address security breaches by maintaining detailed logs for auditing.
- Maintain compliance with local regulations and privacy laws that are applicable to you.

Case study

Integrating AI into ProjectABZ: A project management tool created by ABCDZCo

ABCDZCo is an established and trusted traditional software company that's built a brand identity and reputation for building effective operational tools over the course of the past 22 years. Their overarching company values are productivity, efficacy, trust, and innovation. Their flagship product *ProjectABZ*, a project management tool, has been a dependable staple for small to medium-sized businesses since 2015. They've amassed a loyal customer base due to their unmatched customer service, self-service product capabilities, light onboarding, and ease-of-use user experience.

While ABCDZCo has enjoyed years of consistent sales, robust customer lifetime value metrics, and positive customer feedback, they know they can move the needle further by bringing their product into the future with AI capabilities and features. This poses a double-edged sword for ABCDZCo. On the one hand, they have a soft landing with which to test out new AI features with customers that they've consistently built trust with for years. On the other, executing an AI strategy and design incorrectly could impact the good graces they've painstakingly built over time.

Despite the risks, ABCDZCo is ready to take on the journey of building AI enablement into their ProjectABZ product design, and they're willing to roll the dice. Their plan is to test out the AI capabilities and features they decide on with a close-knit group of trusted super-users and beta testers. They've also isolated a key number of customers that will represent the initial phase of their AI feature rollout before altering their product experience for all customers. While this adds more time to their design and development process, they know this will hedge them against the adverse effects of AI adoption gone wrong. The upside of winning with customers, enhancing their product functionality in a meaningful way, and staying competitive in their evolving market means the potential gains are well beyond their potential losses. Let's look at the different stages of the product development lifecycle of ProjectABZ.

Ideation and research

The product team at ProjectABZ had a clear objective when revisiting their original product ideations and research analysis: they wanted to integrate AI in a way that substantially improves the productivity, user experience, and analytical capabilities of their product in a way that doesn't compromise their users' trust and security. They understood that this was a big undertaking and made sure to work with leadership to manage their expectations about what this process would look like and how long it would last, as any product team should.

Because ideation can have such a huge impact on your product's trajectory and utility, they made sure to involve stakeholders early and dedicated a significant amount of time to brainstorming sessions. By creating a culture that's inclusive of all viewpoints, the ProjectABZ team made sure that project managers, customer success, developers, data scientists, and leadership were all well represented at these brainstorming sessions. They kept the focus on building a map of the user experience journey to better identify pain points and opportunities where AI implementation could add the most value for ProjectABZ.

Significant time went into surveys and interviews with current ProjectABZ customers and users to understand their needs and concerns with regard to AI implementation, which yielded the following insights:

- Users consistently voiced concerns about ProjectABZ's ability to help them with task management, something that was often a source of feeling overwhelmed.
- They also expressed a heavy demand for forecasting and predictive features to help with planning and budgeting.
- There was concern about the privacy of their data and whether they could trust the determinants of AI systems as a whole.

The ProjectABZ product team didn't want to create design silos. They didn't wait for the full list of potential ideas before running a feasibility study on them. They created a staggered approach where as soon as one idea was identified and found worthy of being on the list, they tested the feasibility of that idea early. While this might sound chaotic to some, doing this allowed them to evaluate the technical feasibility, market appetite, and philosophical alignment of that idea with their overarching company values. Their goal was to have a shortlist of AI features and capabilities that could be developed within the next eight months that also had compelling benefits for their users and customers.

After several rounds of ideation exercises, brainstorming sessions, and feasibility testing sessions, they came up with the following list of top AI implementation opportunities for ProjectABZ:

- Task prioritization
- Predictive project timelines
- Intelligent chatbot
- Automated resource allocation

Let's look at how they went about designing these features.

Design and development

ABCDZCo invested heavily in setting up their data science and ML teams for success by having them start their technical research early in the ideation phase. Just as the product capabilities, use cases, and user experience had to be re-examined, the AI enablement team needed time to research which models, technologies, and frameworks would work best for the AI program they were building. They selected the ML algorithms to be used for task prioritization, resource allocation, and project timeline prediction features, and they used NLP for the chatbot.

They wanted, first and foremost, to design wireframes and prototypes of the AI features in action through a **user-centric design**. This meant there were principles of *transparency*, *control*, and *privacy* that they wanted to build into the product experience from early on.

Managing expectations and offering clear explanations for how certain AI features work, offering onboarding modules for super-users and customizable AI settings for those that were more curious or technologically inclined, were all strategic product design choices that were made early on.

They were always an agile development company and this didn't change when they started to incorporate AI. Remaining agile meant that they could develop and ship product capabilities incrementally and fully deploy those capabilities to all users in a way that trickled through their existing customer base. Here is a summary of how these principles were incorporated into each feature:

- **Task Prioritization**: Users had the option to enable AI to analyze their data at the project level to suggest and recommend a priority list of tasks based on their established deadlines, resources, and dependencies. Trust was conveyed through the explanations provided via an indicator icon for why each suggestion was made. Users had an option to disable or enable explanations as well.
- **Predictive Project Timelines**: Because they were using AI to analyze historical data over time, this feature allowed users to learn from past projects and delivery schedules to better predict when upcoming projects could reasonably be completed by. This feature also had the capability of alerting users to potential delays or risks that could impact a project moving forward or prevent it from completing on time. Users could also have the option of incorporating more aggressive or lenient prediction modeling strategies, allowing for more AI transparency.
- **Intelligent Chatbot**: Offering 24/7 customer support through an NLP-powered chatbot gave ABCDZCo and ProjectABZ the ability to retain customer service as one of their top market differentiators. The intelligent chatbot allowed them to learn about what AI issues customers were most impacted by for better or worse. It also helped them build better-prepared customer support teams for the top issues facing their customers. Because the chatbot was continuously learning and improving, this feature had a positive customer impact but it also positively impacted company operations as well.
- **Automated Resource Allocation**: These AI recommendations came from many of the same data tables as the timeline prediction feature but, instead, focused on the optimal resource allocation strategy for that customer based on the availability of their team and their project requirements. Because users had the ability to control thresholds and override AI suggestions, ProjectABZ ensured users still felt they had the upper hand on this feature.

Marketing and communication

The ProjectABZ product team invested in an in-depth competitive analysis early to better understand how some of their biggest competitors were integrating AI. They noticed many used vague language and a lack of transparency as a way to keep their intellectual property out of the hands of prying competitors. ABCDZCo saw this as an opportunity to differentiate their product and offer a premium experience while also putting their customers' and users' fears about privacy and trust at bay. The overarching decision to communicate and market the new AI features and capabilities of ProjectABZ in a way that emphasized their commitment to trust and transparency not only aligned with ABCDZCo's company values but also served as a stark differentiator for ProjectABZ.

Because of this, ABCDZCo's marketing team saw an opportunity to develop a campaign that highlighted the new AI features while balancing and addressing trust and security concerns about AI head-on. What this looked like for them was a series of webinars, demonstrations, case studies, blog posts, FAQs, and customer testimonials to explain the new AI features and capabilities, data privacy, AI bias measures, and user control options. They posted this content publicly and often to capture their existing user base, as well as to attract newcomers to the project management tool space.

They kept close communication with the beta-testing team of users and customers to ensure their feedback loops came from committed partners they could trust. Super-users were rewarded with free premium features in exchange for their feedback and trust. All feedback was incorporated to help with feature refinement and to address any major product concerns and issues before launch.

Their launch campaign included personalized emails to existing customers and users, as well as prospective customers and users who had already expressed interest in the AI features through engagement with social media or through the beta tester program. Social content was centered on the value the new AI capabilities had on the product and on real customer testimonials. Influencer engagement and partnerships also helped ProjectABZ gain traction with known people in the industry: tech bloggers and industry thought-leaders who were able to review and promote ProjectABZ 2.0.

Summary

In this chapter, we recapped the stages of the AI product development lifecycle and contextualized those stages for products that are going through an AI transformation. We went over how AI PMs can manage expectations during the ideation phase, how to manage data with strategy and integrity, how to map the user experience journey during the R&D phase, and how to prepare for increased scale and adoption during deployment. We also reviewed some of the major ways AI can have an impact on your product experience through insights, automation, and personalization. We had a discussion on the interplay between product design and marketing and the importance of communicating in a way that's consistent with the new product design that will emerge. Finally, we ended the chapter with an overview of the importance of building AI systems into your legacy infrastructure in a way that's safe, secure, and accountable before exploring many of these themes in our case study.

Now that we've covered the various ways AI implementation impacts the design process of a traditional software product, we can start to explore the ways AI impacts the day-to-day management of a newly evolved AI product. When building natively with AI, as we explored in *Part 2* of this book, you're baking your AI and data strategy into your development process from the beginning. You can revisit *Part 2* to understand that more fully. In this scenario in *Part 3*, we're exploring the nuances and logistics of incorporating various aspects related to AI enablement after you've already launched and found success with a traditional software product.

In the next chapter, we'll explore how AI enablement impacts the ephemeral management of an evolved AI product. This chapter will serve as a guide for anyone who's thinking critically about what AI enablement will mean for their product and how they can best prepare for the demands of managing a traditional software product that's presently bolstered by AI capabilities and features.

References

- *What is Salesforce Einstein?*: `https://atrium.ai/resources/what-is-salesforce-einstein-your-2024-guide-to-einstein-ai-products-and-capabilities/`
- *Zendesk's chatbot*: `https://www.zendesk.com/blog/introducing-answer-bot/`
- *Personalized learning paths: Navigating the future of employee training*: `https://www.coursera.org/enterprise/articles/personalized-learning-paths`
- *How IBM Watson overpromised and underdelivered on AI health care*: `https://spectrum.ieee.org/how-ibm-watson-overpromised-and-underdelivered-on-ai-health-care`

16

Managing the Evolving AI Product

In *Part 2* of this book, we covered the aspects of going to market with an AI-native product and the intricacies involved with novice products coming into the market, both from a product design and a product management perspective. In *Part 3*, we've covered the same topics. Only here, it's through the lens of transitioning from a traditional software product to an AI product. This chapter is meant to summarize the elements we've already discussed throughout *Part 3* with a focus on the day-to-day experience of managing an evolved AI product. We will cover the various areas of product management of the evolved AI product and the considerations involved with taking an existing software product that already has:

- An established customer base
- An internal team supporting its development
- Operational infrastructure already set up
- A hospitable market that currently supports an existing product's success

This represents a departure from the considerations we looked at in *Part 2*, which did not build a culture and standard process around the areas highlighted above. This illustrates the point of why AI transformation is so hard for companies to get right. Sure, it's challenging to go to market with a new product. Of course, there are so many considerations to get right when you're launching a new product in a new market. However, re-launching an existing product with AI features and capabilities poses a unique challenge because you already have the infrastructure, cultural norms, and standard procedures set up for the existing version of a traditional software product.

Many things will need to change when you're embarking on AI transformation and evolving your existing product to embrace AI. For starters, the product vision, strategy, and roadmap will all have to be reenvisioned with respect to transitioning to an AI product. This means that, in addition to launching a product in the market and ensuring its commercial success, you're also battling the inertia of the way things were done before internally. You don't have to necessarily worry about this when you're building natively with AI; everything is new there. However, that's not the case when you're re-launching an existing product.

In recognition of that, we will be focusing this chapter on the following areas:

- Reimagining alignment with strategic elements starting from the top: product vision, product strategy, and the product roadmap, as well as their alignments to the overarching company goals and mission
- Establishing a feedback loop that will impact how this evolved product and company vision will be communicated externally and internally
- Establishing trust with your internal stakeholders so that they understand how AI transformation will impact their role and performance
- Fostering a culture of psychological safety and experimentation when working with AI
- Managing the ongoing needs and requirements of supporting an AI program that powers your product from a data, collaboration, training, testing, and feedback perspective
- Mitigating workflow changes and system integrations that can disrupt your ongoing operations

We will close this chapter with another look at our case study example: ProjectABZ made by ABCDZ-Co. This chapter is meant to condense and prioritize many of the areas we covered in *Part 2* in a way that's more accessible for those working in environments where a product has already been stable and successful but is ready to re-launch with new AI capabilities.

The head – managing alignment

We discussed the importance of maintaining product alignment in *Chapter 11*, where we dove deeper into managing the AI-native product. You can refer back to that chapter as many of those principles will be applicable here. In this section, we will examine the importance of alignment and its various facets through the lens of evolving a traditional software product into an AI product. Understanding what product alignment is and how it impacts your business and your new product's positioning is crucial for setting the stage properly for the massive next step of integrating AI.

Managing and maintaining product alignment is the work of ensuring every aspect of a product is in harmony with the established product vision. In the context of transitioning traditional software products into AI products, this means evolving everything, from how the product is developed, marketed, sold, and positioned to how it's engaging with your customers and users, in accordance with the new product vision. This new product vision is also in alignment with the company's overarching vision as well. Because this chapter is focused on the evolved AI product, we will go over the elements of strategic alignment and establish feedback loops and mechanisms that relate to the AI-powered re-launched product. Making sure AI transformation is effective top-down will result in a transition that's as smooth and as successful as possible.

Strategic alignment

Strategic alignment is first and foremost a confluence of four major areas that are important to a product's success:

- Reflection of a business's goals and values
- Demands of the market

- Operational capability
- Meeting customer expectations

When you're re-launching an existing product with new AI capabilities and features, these are the areas you'll be revisiting. As we mentioned prior, it's important to revisit the product vision and goals when evolving products because you want to make sure that *why* you're building hasn't changed. So many factors go into AI transformations. Perhaps it is market influences that are prompting the evolution. Perhaps your CEO got inspired at an AI conference. Perhaps your customers are asking for certain efficiencies or automation. There has to be recognition on the part of the PM that incorporating AI could change the product experience so drastically that it should effectively be thought of as an entirely new product. It's also entirely possible that the problems your traditional product was solving will disappear with the incorporation of AI. Defining what the purpose of AI transformation is, how AI is addressing gaps in the original product, and how this transformation aligns with the vision and goals of the company will all be important conversations for the product organization to lead.

When you already have an established product that people know and use, you run the risk of going to market with a new and improved product that threatens the success you've already built. The new version might not be worth the hype. It could introduce new problems to your product experience that rub users the wrong way. You're faced with a new challenge: positioning your evolved product as an optimal solution to meet the needs of your evolving market and user base. Chances are, if you're contemplating adding AI to your existing product, your competitors are too. This means that now, your market likely has many new and improved AI-powered competing products to choose from. All competitors influence each other in one way or another, so new innovations result in a changed landscape. Your customers and users are responding to these new choices by re-examining and refining their expectations of products to choose from as well.

Even internally, various employees in various job families might have different expectations of what AI transformation will mean for the product they support, and how exactly it will make the lives of your customers and users better downstream. So, the work of maintaining alignment is really about making sure that everyone is on board with the new value proposition of the product you're re-launching. Strategic alignment will profoundly influence how the market will accept your AI-evolved product. In turn, acceptance, adoption, and met and exceeded expectations will result in the long-term success of the new AI-evolved product.

The ongoing work of managing alignment is integral to the work of an AI PM because, in this role, you're constantly keeping an eye on your competitive landscape and your product's standing in that landscape. It's your job to see how AI technology can enhance your product experience and how to best communicate that strength to stakeholders internally and externally. In *Chapter 15*, we discussed the concept of going back to the drawing board; the idea of designing a new product experience that poses a significant improvement to your product experience to increase the chances that it will be a formidable player in your market.

In this chapter, the focus is on the ongoing work. This is an important area to stress here: maintaining alignment is an ongoing endeavor. Sure, you'll spend a significant amount of time designing a product experience that satiates customer pain points and effectively competes in the market. It will be a top-heavy exercise. A lot of this work will happen when you do go back to the drawing board. But it doesn't stop once you have a list of AI features. It will be a constant loop for an AI PM. Falling out of alignment is an issue of entropy. All organizing efforts on your part from the design phase go out the window as an AI PM if you can't maintain strategic alignment on an ongoing basis.

Here are a few suggestions to ensure you have strategic alignment on an ongoing basis:

- A quarterly in-depth market research workshop with stakeholders present to understand and communicate how the needs, preferences, and pain points of your target customers and users are evolving
- A monthly meeting to identify market gaps and opportunities and how they can actively be addressed with AI
- Communication of new features with each significant sprint release and a clear definition of how these enhancements contribute to the value proposition of your product to both customers and users, as well as to all internal stakeholders
- The development of marketing collateral and a cohesive messaging strategy that reflects the consistency between your brand's overall identity and the new AI product's value proposition
- Management of strong feedback loops to make sure you're hearing from direct users often and the integration of that feedback into your product planning and roadmap cycles

No one at the company should be confused about any AI feature's contribution to your value proposition. If your new product's AI capabilities and features are only deeply understood by your leadership, product, and development teams, it means you're not communicating enough. All internal teams should understand the product's goals and positioning and they should be able to contextualize the significance of the AI capabilities you're enhancing your product with and how they bolster your product's value.

Feedback loops

Because of the risks associated with re-releasing an existing product with new AI features and capabilities, securing proper feedback loops is important to mitigate those risks as much as possible.

When we talk about feedback loops, most people imagine user feedback. While user feedback is important, it's only one part of the feedback we need to manage product alignment. Market research, product positioning, and internal stakeholder feedback are also crucial to understanding the health of your product offering.

The following is a list of recommended feedback loops and best practices to ensure you're maintaining proper product alignment on an ongoing basis:

- **User feedback**: Whether you're achieving this through ongoing formal **user experience** (**UX**) research surveys and user interviews or informal in-app surveys and polls, your goal is to amass user opinions on the AI features you've added and their effectiveness and to analyze those findings in a way that's meaningful. Finding ways to incorporate feedback in small ways like rating features or reporting issues directly in your platform with minimal keystrokes also goes a long way to serve as an early signal and pulse check on how your new features are being received by your users. Consider your user feedback data as the tip of the iceberg. Often, users may not have the exact words to articulate what they truly need and feel. So, when you do get any feedback, know that what is being communicated at all is significant.
- **Analytics**: If you can set up product analytics to understand how users are behaving in-app, you can make determinations without necessarily needing to engage directly with users. This shouldn't replace direct user feedback, but it should be a complimentary feedback mechanism. If your analytics can corroborate the story coming directly from your users, you have proof you can rely on for years to come. Monitoring behavioral analytics and feature adoption rates will allow you to see which AI features are being celebrated and which are being ignored so that you can make more informed decisions later. Analytics help you understand which features and capabilities are truly being valued the most because they're based on actual behavior in the app.
- **Support data**: Particularly early on in post-launch, you should keep a strong focus on complaints and issues being flagged up to support. Customer service interactions and support tickets show you where your product experience is most confusing. If a user is taking time out of their day to write a ticket or call your customer service number, that's significant. As a user, I personally avoid having to talk to anyone if I don't have to. It's unpleasant to have to face drawbacks and limitations in your product, but the earlier you can catch them, the sooner you can limit the risks of an AI feature gone wrong. As an AI PM, you won't be responsible for responding to support tickets, but you should take an interest in understanding what issues your users are consistently coming back to complain about the most. Analyzing the content from support tickets and customer service interactions is a good way to find common issues and concerns that have been introduced into your product experience as a result of the AI transformation. This then gives you an opportunity to address these concerns through marketing collateral, thought pieces, or, at the very least, an update to your product's FAQs or knowledge base.
- **Competitive feedback**: Staying informed on how your industry is evolving with AI is important for maintaining market alignment. Conduct market research regularly to understand what features your competitors are adding to their products. Make sure your ongoing AI features are competitive and relevant. This doesn't mean you have to do what your competition is doing. It means you're remaining aware of how their actions compete with yours. If you can position your product's AI features and performance against others in your competitive landscape through the use of benchmarking standards, even better.

- **Internal feedback:** You should be valuing the opinions and needs of your internal stakeholders as well. The AI PM role is cross-functional in nature. You will need to bring in a variety of voices and you should be giving them a platform to air their AI concerns, ideally in the form of a cross-functional team meeting or review of some sort. The discussion of AI feature performance and alignment should not be done in a silo. These meetings should always be collaborative because you need the voices of your data science, engineering, customer support, analytics, product, and leadership teams all present to make sure you're maintaining proper product alignment. Each stakeholder represents expertise in their given domain and how it impacts customers, so all perspectives are valid and valuable.

Remember that this can look like a lot but it doesn't have to be. Setting time in your calendar to make sure all these aspects are well represented and that you're checking in with them on a regular basis is enough. Even if each of these areas can be touched on a bi-monthly basis, that's still often enough that you will keep the feedback loop alive. If too much time passes in between, you run the risk of not catching dynamics that could impact your product experience fast enough. But it's still better than overlooking them completely. Just make sure that whatever cadence you implement feels natural and sustainable for you.

The heart – managing the people and values

People don't do their best work when they're undervalued, underappreciated, or overlooked. Whether we realize it or not, the integration of AI into our workflows and routines can elicit a fear response in many people because of the way AI is being weaponized against human productivity. Bolstering your product experience with AI and managing proper expectations with the humans who build and support your product sends the message that the company is ready to evolve and that it's bringing everyone with it.

Change is hard for people and the adoption of AI represents a significant step that needs to be properly communicated. It's an opportunity to strengthen your connection to the employees at the company in a way that feels supportive and safe. It's also a chance for the organization to help prepare internal teams for what AI integration will mean for them and the work they already do. This is particularly true for AI adoption because most people may not fully appreciate how much of a change this might pose for their daily workflows. Preparing internal stakeholders for how AI adoption will impact users' experience of the product and their own workflows, and how to talk publicly about these changes, will largely fall on the AI PM.

On a foundational level, all internal stakeholders and their teams will have to be well versed in understanding the basics of the AI product development lifecycle. Understanding the basics of AI, ML, and data science will be instrumental in giving everyone a baseline of knowledge to work up from. All employees at a company won't need to understand AI in any great depth. But if your role supports an AI feature or if your role is outward-facing to the point where you're talking to prospective customers and users about an AI feature, you should have foundational knowledge. I've been on calls with customer service and sales folks where it's become abundantly clear that the person on the other end has no idea what AI is and how it works.

It undermines the credibility of the company if your customer-facing representatives can't properly articulate how an AI feature works, and they should all have an understanding of the AI fundamentals and the AI model lifecycle. Anyone who is working in service to your new, evolved AI-powered product should understand the impact of the following areas on your AI product:

- Data collection, data quality, and preprocessing
- What it takes to train, validate, and evaluate model performance
- Relevant metrics and how they contribute to business goals
- How AI features are being deployed and which strategies work best
- Importance of accurate, complete, and consistent data
- Impact of regulation compliance and which laws impact their product
- Which data security measures are most meaningful to their product
- Ensuring stakeholder teams are interdisciplinary and represent a diverse set of skills
- Knowledge of their own roles and responsibilities when it comes to managing their part of the AI product integration process
- A curious culture of experimentation with AI technologies and ideas
- The concept of drift and how models and data samples are subjected to it

Psychological safety is about feeling empowered to bring up an area you don't know much about and to be honest about that. Pushing for internal teams to be encouraged to continuously learn about AI is a great way to help people feel supported by what AI transformation will mean for them. Whether they're in marketing, sales, customer support, or finance, they're all still working in service to an AI product. Encouraging people within your product organization and beyond to expand their education of AI concepts through courses, workshops, certifications, or books like this one will help make information more accessible and will help more people within your organization keep up with the rapid pace of AI.

Unifying stakeholders around a shared goal with as much support as possible means everyone will work toward the same outcome regarding AI transformation. It will also help them align their own workflows and metrics in a way that incentivizes AI success. The more stakeholders understand the major areas that surround the maintenance of an AI system, and how that AI system empowers the AI features you're releasing in your product, the easier it will be to champion your AI features and ideas. If your goal is to have your major stakeholders play an active role in helping you build and support your ongoing product roadmap, along with AI features that you'll be releasing for years to come, you have to make sure they're on board with AI.

The new features and capabilities you're adding will have a direct impact on how you will create the infrastructure to support AI. However, there will be adjustments that all stakeholders will have to get used to in varying degrees. There will be new workflows for data management, model training, and deployment at the very least. Beyond just learning about AI, members of the product team and other stakeholders will need to understand how to work with these new workflows and processes. They may not be aware that this will require continuous iteration and model updates or will impact existing development cycles. They may not even be aware that competing products already have AI features implemented.

AI and data literacy go a long way to dispelling fears around AI and job displacement because as time goes on, they will encounter more and more products that are AI-powered. They will also start to incorporate more AI products into their own workflows irrespective of the AI features you're adding to your product. Because AI transformation is something that will impact most products at some point, you never know where the idea of a new feature may come from. Helping to bridge the gap with regard to AI literacy also allows for literacy in other areas, like the importance of data quality.

The move toward AI transformation and adoption should be communicated as an opportunity internally and externally because it is. Externally, it represents an evolution of what your product offering can do for your users and market. Internally, it represents a chance for internal teams to evolve to meet the new product. This will have repercussions that most people can't see yet, so it's best to keep the focus on the positive. Try to limit the negatives as much as possible. Yes, AI transformation is a huge undertaking that will require a significant amount of time and energy to implement successfully, culturally, and functionally. But the more this can be done with joy and safety, the better the whole organization will be for it. The following is a list of best practices and suggestions for handling AI transformation best with your most prized resource – the humans that build and support your product:

- Build a cultural strategy around AI literacy that starts with an AI town hall to get a sense of how internal stakeholders might be threatened by AI and what it represents. Topics can include what AI means to the company, how they're planning to work with it, and interest areas for how AI may transform the existing products the company develops.
- Then, build a series of AI workshops and trainings around those issues and concerns. This is meant to symbolize a cultural shift to embrace AI so these sessions should be handled by a variety of leaders. The burden of communicating and evangelizing AI should never be left to the AI PM alone.
- Ensure teams have a list of external resources where they can bolster their own knowledge beyond that.
- Create a process where stakeholders can communicate interest in broadening their skill set or job description to advance toward AI.
- Identify skill gaps to understand which roles on which teams need to have AI skills and implement a training program to get them there.
- Educate teams on the impact bias, drift, and fairness have on AI and machine learning so that everyone at the company can contribute to an ethical AI product.
- Educate teams on AI system security and potential forms of adversarial attacks on AI models.

The people management side of the PM may not be a fixed line. Many PMs serve as individual contributors in an organization. However, culturally, you will serve as the nexus of AI and how it will impact the product; so, whether you want the role or not, you will be the AI guide for many people in your organization. Partly, this is because much of the communication around what's going on will fall on you. But perhaps more importantly: the more people you have around you who are educated about what AI transformation is and what it will mean for their work, the better off you'll be. It will ensure you're keeping the right perspectives close and it will be easier to manage the highs and lows of executing your AI-driven roadmap downstream.

At the end of the day, AI transformation is really change management. You will need to help prepare your organization for the change that is coming. This will have repercussions that will impact communication plans, training, support, adoption, and political buy-in. Keeping stakeholders informed about the AI integration process, progress, challenges, and successes will engage your stakeholders throughout the process in a way that's supportive of you and your product development.

The guts – managing data, infrastructure, and ongoing maintenance

Now that we've covered product and people alignment at a high level, we'll examine the various areas that will impact the work of an AI PM once you're in the flow of managing your product on an ongoing basis. Understanding the ongoing technical needs, compliance, and risk management involved with an AI product will help curb surprises that may arise downstream, particularly if your stakeholders aren't aware of what AI transformation will mean.

Infrastructure and data

The most obvious area that will change will relate to the infrastructural and data requirements to support an AI system. In many cases, you'll be using historical data to help make projections and forecasts, optimizing current data to offer ranks and recommendations, or creating profiles to identify anomalous behavior. All the models that will be used to support these activities will have to be trained, validated, tested, and retrained on a recurring basis. That means your **computational resources** will be stretched to accommodate the realities of managing an AI system and data storage needs. If all goes well, you scale and your infrastructure will be tested again to accommodate an increase in data volumes and model complexities.

You'll also need to create a plan around how you're going to access data on a recurring basis and establish standards for **data hygiene and pre-processing**. New data pipelines will need to be created to serve your AI system. Various team members and stakeholders will need to be able to have secure access to the data they will need to complete their work. It's likely you'll need to set up new data warehouses and data lakes to support the ongoing AI features and capabilities you're adding to your existing product.

As you mature and scale, you'll benefit from establishing **data governance** practices to make sure you're keeping up with regulations and governance policies. If you don't have an established data ownership strategy, you might want to set one up so that there's a clearer delineation between which teams own which datasets to ensure there's consistent data integrity. This will be especially helpful if you do run into data quality issues downstream, as you'll know exactly where to go to troubleshoot them. You'll also likely want to champion a searchable data catalog tool where you can explore data lineage, particularly as your organization becomes more culturally mature.

Because you're supporting the evolution of an existing product, this means you'll need to create an AI system around infrastructure that already exists as well. This means that one of the unique challenges will be understanding how to integrate your AI system into the current software, infrastructure, tools, and data pipelines you've already built so far. Finding **compatibility** between the new features you want to add and the current systems in a way that avoids or limits disruptions to your existing product experience will also be a challenge to the AI PM who's evolving it into a traditional software product. This process of redesigning your product from a technical perspective to support your AI system will result in new workflows and automation to help you manage that system or, at least, to free up as much time as possible for those working on your product to work on things that matter most. You'll need to think about training team members to work with those new workflows. At the very least, they'll need to train on collaborative tools to help them all manage the work that needs to be done. Technically speaking, there's a lot going on to support the ongoing work of evolving a traditional software product into an AI product.

This work requires a strong partnership with engineering. AI PMs will not be solely responsible for doing this at any company because it's outside of the scope of the role. However, AI PMs will need to be aware that this work needs to be done and should do their part in influencing and encouraging progress toward this work. While engineers will ultimately be doing all this work, it's up to AI PMs to present the problem to be solved and the priorities of the work that needs to be done. AI PMs will organize and present the findings to all relevant stakeholders and leadership. The following are a few best practices you can use to better prepare your organization for the data and infrastructural needs of setting up an AI system:

- Conduct a current state analysis to understand what limits and capabilities your current, established infrastructure has. The less you have to change initially, the better. Make efficiencies where you can.
- Test your current infrastructure for scalability, particularly for computational and storage demands for supporting your proposed AI features.
- Inventory your current data sources, tables, transformations, and pipelines.
- Gather requirements or the data needs your proposed AI features will have and confirm that you have all the data you'll need to support those feature requirements. If you have gaps in your data, create a plan for how you'll fill them.
- Leverage cloud services that are adjustable and flexible wherever you can. Most cloud services are charging you for real usage and this can be advantageous early on to help with fluctuations in storage, AI microservices, and processing power needs.
- Build robust data pipelines for data collection, processing, storage, and mining.

Being thoughtful about the data and infrastructural needs early on will make it so that you can better anticipate and maintain the needs of your AI system as time goes on. Remember that you'll always be making incremental improvements. The work of maintaining your product and the AI system that powers it will be an ongoing effort. Particularly as you scale, you'll be making necessary improvements to make sure your product is working as intended. Next, we'll look at various aspects of maintenance that will be supplemental to the work you're already doing in managing a traditional software product.

Maintenance

On top of the work you're doing to manage a traditional software product and make sure the product or platform is running smoothly, you're also now going to be managing an AI system to support it. Setting up systems to monitor AI models, their performance in production, version control, and tracking the results and metrics of that model performance will be an ongoing effort that needs to be regularly checked. For resource planning, you'll want to keep track of the AI system's resource usage to keep track of what it's costing. But you'll also want to be tracking accuracy, latency, and other metrics to make sure you're within the threshold that's accessible for maintaining a strong customer experience. Support and maintenance are key requirements that are often overlooked and minimized by key stakeholders and leadership. AI PMs will need to present this dependency as a critical requirement and prioritize tasks and roadmaps accordingly to reflect the importance of maintenance.

Much of the work of maintaining an AI system will be for the goal of keeping up with this threshold you've set for acceptable **model performance**. In addition to the regular system maintenance you're used to in traditional software development, you'll also have to set up a regular cadence to retrain your models with new data to make sure they keep performing in a way that's accurate and relevant. Any deviation from that should have a trigger that alerts you about an anomaly. Anomalous activity could indicate model drift, so looking out for those occurrences and setting up guardrails in place will ensure your models are operating optimally.

Beyond logs and performance tracking for the models, you'll also be maintaining **product performance metrics** and KPIs. Though you'll already be tracking many of these in an existing traditional software product, you will be adding new AI-specific metrics to drive your product optimization. Because these are new AI features you're introducing, you'll also want to tie your product performance KPIs to real business outcomes. Finding a way to measure how your AI features are contributing to customer satisfaction, revenue growth, churn, and other improvements in efficiency will inspire your leadership team and give all your peers and stakeholders at the company more confidence in the work you're doing.

Even **documentation** will be an active maintenance. All AI processes, models, workflows, and procedures need to be documented so that all relevant members have access to the optimal flow. As that flow changes and gets improved over time, your documentation efforts will need to reflect the current standing of the systems that support your product today. If knowledge is centralized around a particular team and individual and the information isn't documented, you run the risk of a significant loss to the business should something happen to those individuals.

Overall, the emergent AI system will pose more of a threat and a risk to the business. There are significant risks that come from a lack of knowledge transfer and documentation. There are even bigger risks that come from a lack of security around data and the AI system as a whole. Making sure data is encrypted both when it's at rest and in transit when using APIs and giving specific team members access to specific data samples or AI models will limit the exposure of your data to unauthorized actors. Doing regular vulnerability assessments and penetration tests to address where your AI systems might be weakest will be an ongoing security effort as well.

Case study

In *Chapter 15*, we introduced the case study example of ABCDZCo, a software company that evolved its project management product, ProjectABZ, with new AI features and capabilities. We covered the various areas in which AI impacted ProjectABZ from a product design perspective pre-launch. In this chapter, we will revisit the same case study example. This time, we will go over the various areas AI has impacted ProjectABZ from a day-to-day product management perspective post-launch.

To refresh our memory: the goal of evolving ProjectABZ with AI was to first and foremost enhance the experience their customers and users had of their product. With AI, they were also able to improve the efficiency of their product and gain competitive strength in their market. For convenience, the various AI features and capabilities for ProjectABZ resulted in the following enhancements to the product:

- Task prioritization
- Predictive project timelines
- Intelligent chatbot
- Automated resource allocation

First, let's explore what AI transformation looks like for ProjectABZ before we explore the various ways ABCDZCo managed alignment across management, stakeholder teams, and operations.

AI transformation for ProjectABZ

In this section, we will cover how ProjectABZ changed from an ongoing product management perspective and we will also discuss how this impacted product alignment, people management, and operational management post re-launch. After a transitional period of trial and error, there were several adjustments that were introduced to their product management experience that were most notable post-launch:

- The product team shifted to supporting daily product stand-ups. Many development teams already do this, irrespective of whether a product is involved or not. The ProjectABZ team found that having a daily stand-up for the pre-launch and post-launch periods was instrumental in ensuring product alignment from a development perspective.

  These stand-ups were not just for developers and PMs. They were cross-functionally represented. They had data scientists, AI engineers, and product managers actively participating in these meetings with the option for these roles to log off after the first 30 minutes. This ensured that development stakeholders were aligned on a daily basis so that all involved could be aware of updates that impacted AI model training, data quality issues, and integration issues.

- The impact of integrating AI into their existing project management product meant that sprint cycles had to get updated. It wasn't enough to split the product development into epics for various product areas of ProjectABZ; they had to now incorporate new epics around areas like data collection, model training, and model iteration. Even post-launch, these epics were static because the work of product managing an AI product is ongoing with constant iterations.

  Before the re-launch, ProjectABZ was a well-oiled machine. It had been released on the market for at least nine years and it was a mature product that didn't need the level of oversight its newly evolved AI-powered counterpart did. Because of this, they went from three-week sprints to weekly sprints to capture all the elements that were ongoing for the re-launch. Having sprints more frequently also allowed them to respond quickly to evolving AI components and the challenges that came up when integrating them into an existing system.

- The last major change for the re-launch surrounded the user feedback they were incorporating. Prior to AI integration, the ProjectABZ product team was able to get user feedback organically since the original launch nine years prior. They had partners, user interviews, and marquee customers who provided feedback on a rolling basis and they were less methodical about their approach. However, they couldn't risk that for this re-launch; too much was at stake.

  The product team at ProjectABZ knew that they had to re-focus their efforts on receiving feedback from surveys, beta testers, and in-depth user interviews. They also added a focus on UX research to manage the ongoing, regular analysis of the feedback they were getting from those groups. This feedback was then frequently discussed with internal stakeholders and teams to refine the AI algorithms further and improve UX post-launch.

Management alignment

After the work done on ProjectABZ (outlined in the *Case study* section in *Chapter 15*), the product team was able to work with their leadership team to update a strategic vision of ProjectABZ that exalted its position in the market as an AI project management tool. They updated their product roadmap to reflect the new AI features, capabilities, enhancements, and integrations that were agreed upon during the design and research phase. Strategic alignment was achieved through in-depth product strategy sessions that explored how the evolved ProjectABZ product experience reflected the company's overarching business goals and values. These sessions were managed by the AI PM and incorporated a lot of the market research and analysis that justified the integration of AI in the first place.

They also incorporated much of the customer feedback and research they did to understand where the product already fell short of expectations. This resulted in a feedback loop that also strengthened their operational readiness and capacity for AI integration. By positioning today's gaps in the product experience as an opportunity for AI integration, the product team was able to reinforce AI transformation for ProjectABZ as a strategically relevant endeavor. They spent a considerable amount of time building their market positioning through marketing campaigns that emphasized the new AI capabilities in a way that differentiated themselves from the competitors in their market. They also made sure these outgoing communications were well understood by internal stakeholder teams to ensure what was being reflected out into the world was also internalized within the organization. They focused their communications on the new value proposition: how the new AI features would reduce project delays and impact productivity for their users.

People alignment

ABCDZCo was already aware of the skills gap and AI literacy requirements they would need to invest in to ensure their organization as a whole was ready for the re-launch. This meant that they developed an extensive training program to help upskill their existing employees to better understand the requirements and demands of AI and machine learning. They hired outside talent for applied AI work like model development and research. However, they recognized they still needed to bolster their organization's overall readiness for AI adoption.

ABCDZCo also understood that their first foray into AI was focused on their product, ProjectABZ, but they had other products they sold. ProjectABZ served as a first use case for AI evolution. For this reason, they were open to piloting a re-training program because they knew they could both retain and upskill talent through the work of investing in their people first. They saw that AI integration would result in a cultural shift and they wanted to get ahead of the AI fear-mongering about job loss by addressing it head-on, reinforcing that no one was losing their jobs. If anything, they were going to be receiving the benefits of AI productivity gains and new skills upgrading.

Just the idea of AI led to a cultural revolution within ABCDZCo, independent of the ProjectABZ product team, in which a culture of continuous learning and innovation was created and grown within ABCDZCo. This led to them becoming more desirable from a talent perspective as well. Training and workshops were developed to cover applied AI fundamentals, AI ethics, and bias mitigation, as well as focused workshops for marketing, sales, and operations.

Finally, the AI re-launch actually led to more interdisciplinary, cross-functional collaboration company-wide. The need for more collaboration due to AI resulted in additional benefits that broke down the silos that existed between development and operations teams prior to AI. Because so many were upskilling and collaborating more, this meant that knowledge-sharing sessions were becoming more common for both the ProjectABZ product team and other teams that supported ABCDZCo products. AI innovations, breakthroughs, developments, and best practices were routinely shared between stakeholder teams as a result.

Operational alignment

From an operational perspective, both leaders at the ABCDZCo company level and the ProjectABZ product level understood that an AI re-launch was going to mean a significant financial investment. Computational resources, cloud services, AI processing hardware, and increased AI-focused head count all contributed to the costs of maintaining a pilot AI program. However, setting up the proper infrastructure to support ProjectABZ meant that ABCDZCo could then use that same infrastructure to support their other products as they scaled and increased their margins with ProjectABZ. Data pipelines also had to be updated and optimized to make sure the flow of high-quality, clean data would feed their AI models.

From a maintenance perspective, new workflows had to be created to handle the new supply of data feeding the models. AI model training, validation, deployment, and continuous monitoring had to have established workflows of their own. The team also had to build out automation to handle repetitive tasks and set up flags and logs to oversee all the new workflows and activities that had to be created to support and monitor the new AI program, particularly as it related to AI model performance, anomaly detection, and troubleshooting. This meant that there had to be regular model retraining schedules with the latest, refreshed data, as well as security protocols to strengthen and protect their data and AI models from breaches and adversarial attacks.

Results and outcomes

Nine months post-re-launch, this resulted in the following product outcomes for ProjectABZ:

- **Efficiency improved**: The evolved product led to a 43% reduction in project delays and efficient task management. The automated resource allocation helped project teams make more informed decisions and an 18% increase in project success rates.
- **UX improved**: In-product summaries reflected a 36% increase in user satisfaction levels compared to the original version of the product with positive feedback surrounding the new AI features, with users finding them intelligent, intuitive, and valuable in helping them with their work. The new chatbot feature also helped users reduce their admin time on the platform by a whopping 63%.
- **Differentiation improved:** ProjectABZ did gain a competitive edge in the market and the new features helped attract customers that previously churned for other competitors. They were able to re-attract 22% of their churned customers from the past three years. The re-launch also helped rebrand ABCDZCo as a future-oriented, forward-thinking company. ABCDZCo also experienced a 45% reduction in churned employees because of how they handled the cultural impact of AI product integration.

Summary

The integration of AI poses a number of challenges for all companies that will face AI transformation. It's not as simple as adding headcount and bolstering infrastructure. It needs to be handled with intention. Many companies will alienate their workforce and their customer base if they don't take on AI transformation with the attention, detail, and integrity it requires. The decision to add AI features and capabilities is one many companies will come to by the end of the decade by necessity. But how prepared they are for it, how they will empower their teams, and how they will execute their AI strategy will make all the difference for employees, peers, customers, users, and critics alike.

In this chapter, we discussed that as the AI landscape continues to change, leadership and product teams will need to remain nimble to the necessities that change will require. We've constantly referenced the iterative nature of AI. We will all be continuously challenged by new players, technological breakthroughs, and infrastructural limits that continue to impact the market. This is why building ethically doesn't just have to do with covering your risk from a bias or data perspective. No doubt these aspects of AI product management are crucial. Building ethically is also about how you'll bring the entire ecosystem that supports your product with you.

As the AI PM, you're going to be the one all stakeholder teams come to for guidance. Whether it's in your job description or not, you'll be one of the resounding voices that will influence how your organization gets AI-ready. Getting AI-ready is about technological, developmental, and operational readiness. But if you ignore the cultural and emotional readiness that needs to come with AI integration, your teams will be influencing your product and building in a culture of fear. Product management has always been about people first. Your role is one where you depend on others to carry out the agreements you make. This doesn't go away when you start playing with AI. If anything, it gets stronger.

This concludes *Part 3* of this book, where we focused on products that are being evolved into AI products. In this section, we looked at some of the considerations companies can take as they're thinking about what AI transformation looks like for them. We looked at trends and insights across industries, how to build a strategy around your new AI-evolved product, and the role of design when products embrace AI, as well as considerations to keep in mind when actively managing your next-generation AI product. The next chapter will begin *Part 4* of our book, where we will take a step back from this applied work to explore the AI PM career and how someone can get started in it.

Join us on Discord

Read this book alongside other users, AI experts, and the author himself.

Ask questions, provide solutions to other readers, chat with the author via Ask Me Anything sessions, and much more.

Scan the QR code or visit the link to join the community.

`https://packt.link/aipm`

Part 4

Managing the AI PM Career

The rising prominence of AI PMs necessitates guiding you through foundational skills and knowledge, strategic insights, ongoing development practices, and continuous learning essentials for those looking to excel in the role.

The fourth part of this book will focus on those looking to start, excel, and mature in their careers as AI PMs. We will focus on how AI PMs can grow, thrive, and evolve into impactful leaders who spark innovation and mature into impactful professionals who contribute value to their organizations and industries. We'll start with a comprehensive introduction for those aiming to launch their careers as AI PMs, continue onto the characteristics and expectations of an effective AI PM, and end with a review of the long-term growth and career evolution of AI PMs that foster a mindset of lifelong learning, openness to feedback, and generosity of spirit.

This part comprises the following chapters:

- *Chapter 17*, *Starting a Career as an AI PM*
- *Chapter 18*, *What Does It Mean to Be a Good AI PM?*
- *Chapter 19*, *The Evergreen Learner: Maturing and growing in AI PM*

17

Starting a Career as an AI PM

It's no secret why the AI PM role has come into focus in recent years. The fast pace of AI, coupled with the budding discipline that is product management, makes the role attractive for a lot of people who are already familiar with tech. *Part 1* of this book focused on the AI landscape to ensure PMs coming into the space could find accessible concepts surrounding all the various forms of **machine learning** (**ML**) and how they can be employed across products. *Parts 2* and *3* of the book focused on the native AI product and the evolved AI product, respectively. That covers the two potential types of AI PM roles you'll come across. I would consider this book as a very early step for someone considering a career as an AI PM. It's meant to catch the interest of those who have musings about what an AI PM career could look like but haven't taken the first major step yet. This part (*Part 4*) of the book is about crafting what that next step might look like.

In this chapter, we'll explore what attracts people to the AI PM role, which vantage points they largely come from, and key areas in which this career path needs diversity. We will also cover the material and background that best sets up people interested in this line of work. Finally, we'll end with a few of the various paths PMs can take as they branch out into various specializations AI and within product management.

I've personally been in the industry for fifteen years and I had the chance to see the PM role gain prominence shortly after I began my own career in tech. Particularly early on, I remember thinking, "What do PMs actually do?" Some days, it seemed like they did quite a lot. Others it seemed like they were glorified brand evangelists. Often, it didn't feel like it was a job at all.

Over the course of my career, I have come to learn that the PMs who make their roles look easy are the ones who are most engaged. I originally started my career in account management for SaaS B2B companies. In the early days, I didn't work too closely with PMs but because of my proximity to customers and their overarching pain points, I did have the opportunity to collaborate with PMs on occasion. As my account management career progressed, I found myself working with PMs more frequently. I represented the voice of the customer and I had the privilege of working with PMs who appreciated my perspective.

During the course of my career, I have learned that not all PMs are as customer-focused as I am. I've even worked with PMs who considered speaking directly with customers as beneath them. Today, it feels wild that someone in this role might have that perspective; that's a massive gap in their understanding of their job function. Customer experience and customer success are a big focus area for a tech PM because they're disciplines that affect how customers and users actually interact with the product experience you're creating.

But the idea that a PM reading this book might have gaps in their understanding of the AI PM role didn't sit well with me. Gaps like this have the power to make or break your performance in a role or even in an entire career.

Let's take a deep dive to better understand how people enter and stay on this rewarding and challenging career path. This chapter will focus on those looking to start a career as an AI PM, and we will be exploring the following topics as the chapter progresses:

- How to get ready as an AI PM
- The current state of the AI PM landscape
- The importance of maintaining communities
- Interview skills and recommendations
- The paths available to you as a novice AI PM

We won't be spending any time on what the AI PM role looks like here because we've covered the various sides of this throughout this book already. We will instead focus on applied, practical steps to get your career going because, chances are, if you're still interested in it up to this point, you're probably ready to take the next step.

Bolstering your knowledge in theory and practice

A healthy blend of theory and practical experience is essential for building competency. In this section, we will explore an overview of skills, resources, and credentials that can set you up for success when you're approaching a career as an AI PM. Theory provides a solid foundation for understanding AI concepts like ML algorithms, data science, and AI ethics, but practical experience helps an AI PM hone in on real-world challenges like scaling, model deployment, and integrating user feedback. Combining the two allows AI PMs to build a career around making informed product decisions and anticipating potential issues.

Theory

We started this book with an overview of the ML landscape in *Part 1* because it represents the theoretical foundation we need to build on. If you're interested in a career as an AI PM, you'll have to show that you have some understanding of the basic building blocks of AI. You won't be able to make informed decisions or have meaningful technical discussions with your counterparts on ML and data science teams without this. Algorithms, data transformations and processing, model preparation and training, AI bias, and ethics are all important areas of ML that you'll need to understand well in order to have credibility as an AI PM.

There are many learning paths you can take to get there and the option you take is largely based on your own learning preferences. Someone can learn these concepts all for free on YouTube if they want to. If anything, the democratization of concepts and ideas that give people financial and career autonomy is one of the shining pillars of tech's glory. But not everyone will have the patience, context, or drive to keep steady progress alone.

For those who don't, there are a number of online resources to lean on. **Massive online open courses** (**MOOCs**) like Coursera, edX, and Udacity bring a wealth of knowledge for very little financial investment. The financial investment will go up depending on how much support you'll need to be held accountable. MOOCs will have some level of accountability because you have to actually complete the course in order to get the certificate but they're made to be easy to understand and easy to complete as well. A few reputable ones to note are the professional certificates in ML and AI from MIT and the AI Product Management Specialization course from Coursera.

Depending on how deep you want to go with the material, the information there might be superficial in that it's just enough to get you to complete that particular course. I recommend MOOCs for those who are already PMs but want to make the jump into the AI subject matter. The reason for that is that it can be daunting to build up enough credibility through MOOCs alone. People tend to undervalue and undermine the efforts people make on their own.

Without a certificate or degree, there's an underlying assumption that the knowledge acquired on one's own merits isn't substantial enough. But that probably has more to do with gatekeeping and resentment. If you have hiring power over someone, you may be inclined to reward someone who's followed in your own footsteps, particularly if those footsteps were an expensive degree from a US university. But this is not my belief. I believe anyone can get a comprehensive enough understanding of ML to be an AI PM even with completely free resources.

Bootcamps are the next phase up from that because they often come with a mentor, a seasoned professional in the field, who you can go to with all your burning theoretical questions. But be discerning with these. Many offer a money-back guarantee but that's slowly falling out of favor when coming up against an unpredictable job market. Even if they don't offer a job guarantee, they should be able to give you confidence to help you understand, retain, and demonstrate your new ML skills. Typically, these programs come with career coaching as well, so you'll often get a coach who will help you to interview for jobs as a result.

For those who are serious about getting their AI PM career going, I recommend spending at least a year or two in a data science or ML engineer role if you can, particularly if you take the bootcamp or MOOC route. As an AI PM, you're going to be working with data scientists and ML engineers quite often. You'll want to understand the challenges they're up against when they're working in production and it can be hard to empathize with them if you haven't faced these challenges yourself yet.

Can you be an AI PM without this? Absolutely you can. But you will find yourself coming up against perceived and actual limitations in your work, and the softer the cushion you can give yourself to fall back on, the better. Until recently, colleges and universities were the way into most white-collar jobs. Companies wanted to see that you had some kind of pedigree. They wanted to make sure you could show up on time to the same institution for a minimum of four years. No doubt these are admirable skills to demonstrate to any employer, but today, a degree is not the only path forward. This is particularly true if you already have experience in tech or product fields.

For those looking to take the traditional path who have little exposure to the AI or product tech world and want to take their time with the material with little resistance from the outside world: you can always get a degree. A masters in AI or an MBA will serve you well for a career in product management. It's the longer, more time-consuming path, but that comes with its advantages. Being able to contextualize and understand what part of the AI ecosystem most excites you over the course of years is a blessing. When I started interviewing for roles as a data scientist, I encountered a lot of suspicion, particularly from people who hadn't taken that particular route. I knew that if I had an advanced degree, I may have been an easier pill to swallow for some of the prospective roles I was considering.

Remember that learning new disciplines takes time. You're not studying for a standardized test. There is no arbitrary finish line. The field of AI is wide and deep. It takes time to even understand which part of the AI landscape you want to work with. What kind of ML you naturally gravitate to is more of an intuitive exercise than it is a practical one. Depending on your interests and financial resources, go with the option that gives you the most joy.

If you get excited by the idea of learning something on your own and applying it right away, go down the free or mostly free route. Impress your prospective employers with your can-do attitude and your ability to self-start. If you need accountability and more structure, go with a bootcamp. Show off your portfolio and innovative ideas to a start-up that's looking for someone who can learn quickly and apply their knowledge tangibly. If you prefer the pace and depth of a four-year degree, go for that and let employers know that you're in it for the long haul.

The path you take in theory is largely up to you because it's not necessarily your knowledge that matters for building a strong career; it's the practical application of that knowledge that sets you apart from others. So, let's explore how you can demonstrate your AI skills in a way that will help set you up for success in your ongoing career.

Practice

When we're early on in our careers, we spend a considerable amount of time looking at various resources and putting together a portfolio. I entered the AI PM world through my upskilling journey and work as a data scientist. During that time, I spent a lot of energy crafting the perfect portfolio. I thought deeply about the subject areas I would organize my projects around so that I could anticipate a future role that felt right for my interests and talents. It was also a way to show off different skill sets for different disciplines in data science. I recommend taking a similar approach with a product management portfolio.

Being able to prove you can navigate complex challenges and innovate using ML and data science is your task as a novice AI PM. But demonstrating this is not straightforward. In order to build a portfolio as a PM, people have to trust you enough to actually be one. That poses a hurdle for those who have their eyes set on this career path, but it's not impossible.

The following are some options you can use to demonstrate your AI/ML skills:

- Project portfolio
- Competitions
- Open source projects
- Apprenticeships/internships
- Entry-level positions

Particularly for those looking to jump straight into an AI PM role without first working as a data scientist or ML engineer, building a project portfolio of personal projects is a great way to demonstrate your knowledge of AI/ML. You can focus on projects that are more accessible and relevant in today's job market by choosing projects like specialized chatbots or recommendation systems.

Beyond that, you can also flex your tech skills by enlisting in Kaggle competitions. Kaggle is a well-known and reputable community, so you get bragging rights if you can prove that your work was successful enough to win a championship.

AI for Good projects tend to be open source because their appeal resonates with a lot of people at once and many experienced people are often willing to offer their time, expertise, and skills for causes they care about. They're also a noble pursuit no matter what level you're at in your career because we should all be using advanced technology for the common good. If you're early in your ML or product management journey, see if there are open source *AI for Good* projects you can get involved with so that you can demonstrate your skills toward a goal that's mutually beneficial for all. It also increases the likelihood that others will give you guidance or assistance.

The program I did as I was upskilling was called Thinkful and I worked under a senior data scientist as I was building my portfolio. What that meant was I was able to build a portfolio around real work that was coming from the industry. The advantage of this was that instead of doing a project with all the public data samples everyone was using to build a portfolio, I got to focus on real-world case study examples and data samples that came with all the data hygiene issues you're likely to find in production at a real company. It also meant that I got to work under someone who was doing the kind of work I'd be doing five years down the line.

Apprentice programs and internships are a great way to get clout as a data scientist, ML engineer, or product manager because they're able to give you organic, realistic resources to work with that will emulate what you come across as a real PM. They often come with opportunities to work on real problems with the help of a more established working professional, which is preferable to the theoretical knowledge you'd get from courses or MOOCs alone.

I recommend that people try and get some experience under their belt before trying to secure their first real role. But with that said, everyone is different. Grit and positioning play a vital role, and if it's something that's for you, you just might get lucky. Regardless of how you get there, make sure that your first role in the AI PM job family is one where you can receive guidance and direction from others further along than you. I say this very much from personal experience. I was rewarded for my determination and grit, and I got my first PM role while I was searching for a role as a data scientist. Fate took a turn for me and I accepted my first product role as an experiment. I really didn't know whether or not I'd like product management but I decided to try it out and see if it was for me. Not only was product for me but once I took on the role, I felt like all my previous experiences had prepared me for it.

While that's a story with a happy ending, I endured a lot of emotional tribulations. Being thrown into a product role at a small start-up meant that I didn't have the guidance I could have really used early in my product career. I didn't have more advanced practitioners to ask for support or advice. I had to find my way in the dark or ask strangers on LinkedIn. I internalized a lot of things I didn't need to take on. I often felt like an imposter, feeling that I somehow cheated my way into a job I didn't belong in. It didn't matter that I had a working portfolio, several open source projects on my resume, an apprenticeship program, a volunteering history, published articles, or real work experience. If you are able to make a choice for your first product role, I recommend going with the one that will help nurture your confidence and support the best.

The early days of a career as an AI PM are all about demonstrating technical proficiency, business acumen, communication, leadership, and problem-solving skills.

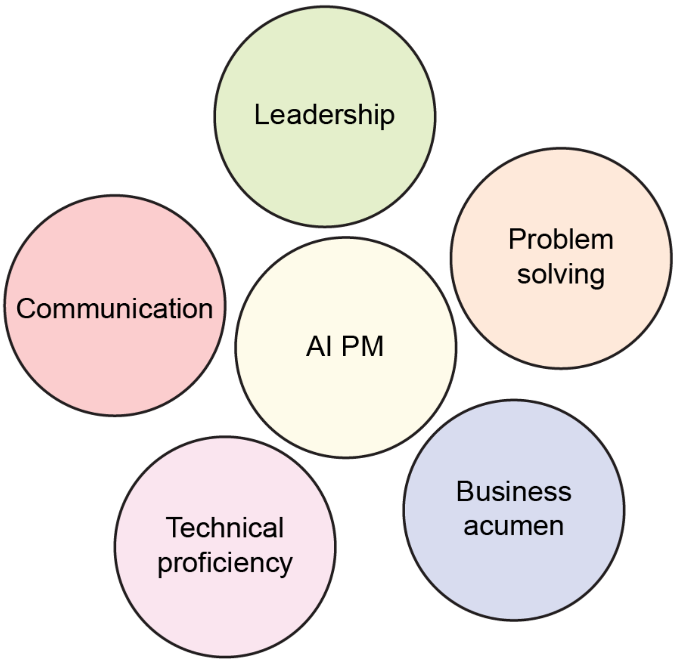

Figure 17.1: Core competencies of an AI PM

We've gone through a few learning and portfolio-building options in this section, but remember that these are the skills you're going to try to convey to land your first couple of roles. Make sure that the work you're doing to learn doesn't just apply to technical AI concepts but that it can be applicable across all five areas.

What an AI PM looks like today

Today's AI PMs are heavy on the business side because they need to possess strong commercialization skills to make their applied AI products a success. Things like understanding market analysis, product strategy, and financial planning are important skills for AI PMs to master because there's simply a lot of market uncertainty. The AI wave is still new and, depending on where you stand, it could feel like a tsunami. We're also early in the hype cycle, so having AI PMs who know how to talk about their brand of AI in a way that's accessible is valuable. Also, because many of the AI use cases are still novel, someone who can translate technical ability into business value and ROI is a huge asset to their organization.

They're not just doing this for their organization's leaders but also functionally for all the stakeholders they work with. We've talked at length throughout the book about how cross-functional and collaborative the AI PM role is, but this translation is highly relevant here as well. AI PMs have to communicate the business value that comes from AI to virtually all their stakeholders on various other teams all the time. With data science and ML stakeholders, they're doing the opposite: they're translating business problems into tech terms. The AI PM has to be bidirectional, effectively communicating and fostering a team environment with every department they work with.

That sounds like a lot because it is. All that communication doesn't happen in a vacuum. There are important projects related to their product with every sprint, so they need to manage the development and deployment of multiple aspects of their AI products and platforms. They're balancing timelines, resources, priorities, and outcomes related to all aspects of their product or platform and they're communicating all that information across teams—occasionally, even to the whole company at once. Depending on where you work, this may not all be happening at the same time by one person. At a small start-up, it likely is. At a bigger start-up, there are anywhere from 5 to 20 PMs. At bigger companies, you have longer, hierarchical product management structures.

But today's AI PM also has a foundation of technical knowledge through computer science, data science, or otherwise related technical fields. Even if they don't possess an advanced degree, they need to have a foundational knowledge of ML to be able to do the job. Chances are, they have had hands-on experience with Python or R and have personally used frameworks like PyTorch or TensorFlow to at the very least understand the full lifecycle of an ML model.

I find it surprising that a technical background is prized well above others in a product role. As I started to reach out to people on LinkedIn to understand their level of satisfaction with their current role as PM, I noticed something many of them had in common: they used to be engineers. This notion is starting to change, but it still holds true that today's PMs are overwhelmingly represented by technical backgrounds. It's a shame really. Engineering skills are great; it's helpful to know about technical feasibility when you're scoping features and technical complexity. But if the bulk of PMs are coming from technical backgrounds, they have a lot of PM skills to acquire if they're going to be successful PMs.

Things like customer empathy, stakeholder communication with non-technical people, project management, market analysis, storytelling, and product marketing are important components for an AI PM.

As time passes and we continue with the AI PM role, particularly as market adoption of AI continues, we will likely see the "AI" moniker drop from the PM descriptor. I don't believe PM roles will disappear over time. Even with more and more "technical PM" roles springing up, we'll likely see that holistic roles devoted completely to the overseeing of a specific subject area of the project stay strong. But over time, it will likely be that all PMs need to have some understanding of AI or ML because it will start to be pervasive. It already is to a large degree because virtually all tech products these days have some AI functionality, capabilities, or features.

The passage of time is likely going to put more of an emphasis on the ethical considerations of AI, particularly as we start to see more regulations, standards, legislation, and legal precedents impacting more consumers and companies. There is a growing importance today for AI PMs to understand ethics and issues of bias for them to do their work well. But this is largely left up to the morality of each individual. Today's PMs want to know if their product is susceptible to model or data drift to the point that it compromises the product experience and introduces bias. They also might be aware of ensuring that they're thinking critically about whether a product is risky or unethical enough that it might attract legal retribution downstream.

But functionally, most tech companies today are not at the frontier of championing AI ethics. Many may pay lip service and craft convincing landing strips about their commitment to building an equitable future with AI. Some may even have entire AI ethics teams set up at their organizations. But when those teams push back and flag an issue that would require a more foundational change, the organization cannibalizes the team. I don't think this is because today's PMs are morally inferior, but simply because there isn't much legal recourse or legislative infrastructure to truly support an ethical deployment of AI.

I'm a firm believer that achieving this more ethical, circumspect future also aligns with an insurgence of non-technical backgrounds in AI PMs. While it's nice that PMs coming from careers as developers, engineers, or programmers can empathize with their counterparts on development teams, it does create a homogenous culture around the work of a PM. But the PM role is much too diverse for that to be helpful. We would benefit from the unique experiences and perspectives of people coming into PM work from humanities, social sciences, and the arts. In particular, as AI use cases and products expand, we would benefit greatly from various perspectives on user behavior, ethical considerations, and even societal outcomes.

In my case, my data science background bolstered my technical acumen but my educational background is in economics and international relations. I think these two areas of study were what prepared me for a career in product the most. My economics background helped me tremendously with the focus on trade-offs, resource management, scope allocation, and market analysis. Likewise, my international relations studies helped me the most with office politics and navigating difficult conversations around prioritization and the building of a strategic roadmap. We don't know how interdisciplinary expertise will come in handy. Demanding roles like that of the AI PM can benefit from well-rounded backgrounds.

Beyond educational and background diversity, we're still starved for demographic diversity in AI product management. Generally, the tech field is still aggressively male-dominated. We've already seen issues with tech products (and physical products) being made by men for men. Increasing the representation of women alone could help product designs and launches be more inclusive and find more compelling market adoption that way. The goal of a PM is to craft product experiences and innovations that reflect the needs of a diverse group of users. Making steps toward bringing more women into AI PM roles is a way of achieving that.

This is also the case with cultural, ethnic, location, and age diversity. If we can find a way to attract more types of people, with more perspectives into global markets or user preferences in a way that's unexpected, it will make for an overall better product experience. Whether we admit it or not, we can all learn from one another. No one group has it all figured out. There's inherent wisdom and knowledge in all demographic groups, and no matter what some people may think, we're equals. Even the way we compete inspires us. An innovation that can come from a competing product, which might be due to the demographic or educational background of their PM, could be the thing that changes the industry as a whole. Little things can change everything at any given moment. This is doubly true when thinking about AI products. AI will continue to shape our future, and the more we can encourage diversity today, the more we can trust that we're building future solutions with significant consequences that can benefit all members of society.

The importance of communities

The fact that there is such a lack of diversity is one of the reasons why we have to rely on building strong communities. You never know how a job will go. Every role you say yes to is a gamble. Everyone has the ability to fall through the cracks, but once you do, is anyone there to catch you? Companies don't operate the way they used to. Companies today have learned that they can hire many employees at once, and conduct massive hiring sprees where their needs can be met in the moment to match the market demand and funding they're experiencing.

But what happens when companies experience periods of downturn? They turn to one of the tried and true tricks of rebalancing their P&L: mass layoffs. Where I live, in NYC, it's a cut-throat environment. The best of the best come here to live, work, and build their dream lives. Those who have strong professional networks can combat the repercussions of layoffs more easily than others who are more siloed, particularly if you're building a career in tech, where layoffs are more common based on the demand for their products. Leaning on established communities is also a way to learn from others with more diverse backgrounds than yours.

The stronger our networks, the greater the opportunities for growth, resources, knowledge, and new jobs. Once you start working in the industry, you can start to specialize in an area of AI or certain industry verticals because they become more familiar. The roles you start to look at might be dictated by the opportunities you take on early in your career. The market overwhelmingly rewards those who specialize because hiring managers are largely looking for a snug fit with their job description.

For instance, if you build a career as a cybersecurity AI PM, you might find yourself taking future jobs in cybersecurity because you have the required number of years of that specialty for the next role. But what happens when you stop being curious about cybersecurity and want to branch out? You may find yourself succumbing to the inertia of the job market because it will reward you for staying in your given area of expertise. So leaning on professional communities is a way to stay employed, but it's also a way to branch out into new areas altogether. Even if you've worked in one area for some time, you could meet others interested in a new specialization through your networks, and together you could work on a project to help break into a new vertical.

As we get started in any particular job function, we also start to streamline the kind of work we do. For instance, if you start to specialize in LLMs and conversational AI products, you'll find that you're deepening your knowledge of LLMs but the knowledge you have of other types of ML models starts to wane. Maintaining involvement in communities is also a good way to have access to others who have a deeper knowledge of other areas of AI as you mature in your AI PM career. Being able to share insights about your work, learn from others, and share best practices and industry knowledge is a way of giving back to the community that's supported you.

Communities are also a great way to learn about new workshops, webinars, or training sessions that are more specialized or relevant to your particular area of work as well. Because the AI PM has to be learning on an ongoing basis, this is an effortless way of keeping your skills sharp and staying professionally social. Keeping up with your skills and staying competitive in your job family is important whether you're on the search for a job or not. In some cases, I've even been able to lean on professional communities when I've looked for beta testers for an app I was building.

The reality is that we never know what direction our career might take in the future. Perhaps some will get lucky and find a long-lasting, satisfying, stable career at a company that will stick with them through the long haul. Most will experience a shifting, volatile market. The exciting thing about this moment is that we're going to see more advancements and innovations continue in the area of AI. Keeping up with new breakthroughs, use cases, and insights isn't only going to help you future-proof your career, but it's also going to allow you to be available when new developments arise that you want to be a part of.

Finally, if you're dead set on building a career as a PM, you're going to very quickly find that it's a thankless, demanding job. But that doesn't mean that it's not rewarding. And who knows, you may even work for an organization that passionately celebrates your contributions. But because the role is so multifaceted, eventually, you're going to need emotional support. We can't be all things all the time. Building a supportive community around you can help you prepare for the emotional struggles of being an AI PM. Building a network of peers you can trust around you can help you share experiences and build coping strategies for when things are going less than stellar. Managing stress, maintaining motivation, and celebrating successes will help make all the difference.

Whether it's emotional support, collaboration, skill building, networking, or knowledge sharing, communities play a vital role when you're starting out your career as an AI PM (and beyond). Beyond this, it also teaches you not to see others as competitors. When you're an active participant in contributing to your communities, you start to see that the whole is greater than the sum of its parts. This makes it harder to be jealous of anyone else's success.

If you do find yourself feeling jealous, you can channel that into something more positive by drawing inspiration from the subject of your jealousy instead. You can ask them for an informational interview to understand how they got to where they are and maybe for some guidance. If they're not open to that, you can learn from their career trajectory on LinkedIn and ask your community for guidance instead.

Either way, your network is priceless. If you build it, they will come when you need them. Leverage LinkedIn and Slack groups for specialized areas of focus. If you do take the path of MOOCs, most of them will have a professional networking group you can join after you achieve your certification. Bootcamps and apprenticeship programs will do this as well. Invest in virtual communities but leverage in-person events and communities where you can. We may be in the age of AI but interpersonal relationships still reign supreme. If your professional area is too niche for there to be a group (an unbelievable predicament), start one.

Here is a list of some communities you can think about joining as you're navigating beginning a career as an AI PM:

- **Data Science Society**: A global community offering data science events, challenges, educational resources, and a network of data scientists and ML professionals
- **Mind the Product**: While this is a broader PM community, they often host events, articles, and forums on AI product management
- **AI4ALL**: A nonprofit focused on expanding access to AI education and programs for those interested in AI ethics and product applications
- **DeepLearning.AI**: Led by Andrew Ng, this community provides courses and forums where users can engage with others interested in AI
- **Women in Machine Learning & Data Science (WiMLDS)**: A community focused on diversity in AI and data science careers with chapters worldwide
- **AIPM**: An online community specialized for AI PMs focusing discussions around AI use cases, product-market fit, ethical considerations, and user-centered design for AI applications and features
- **Women in Data**: An online community for women in data careers that focuses on skills gap training, networking, mentoring, and chapter-based talks and activities in local areas
- **Women in AI**: A hands-on AI community with a range of initiatives and programs devoted to AI applications created by women

Choosing your AI PM specialization

There are a number of paths that are available to you as an AI PM. Often, the main problem is just choosing one. Review *Part 1* of this book when you're deciding on this to better understand which areas of AI resonate with you most. This is a great starting point for many AI PMs because the area of AI you choose to predominantly focus on will be the area of AI you will invest in first. The following are a few questions to get you thinking in the right direction as you're starting to build a career as an AI PM:

- Are you most excited by supervised or unsupervised learning models?
- Do you have a preference for deep learning or traditional ML?

- Are you more interested in developing consumer-facing applications or enterprise solutions?
- Which industries (such as healthcare, finance, retail, or transportation) do you find most intriguing?
- What are your strongest technical skills (ML, natural language processing (NLP), or computer vision)?
- Do you already have experience in any specific industry or domain?
- How important is the mission to you? Do you want to work on AI solutions for good?
- What are your areas of interest outside of work? Do you have any hobbies or c? Creative pursuits?

You don't have to have it all figured out as you're starting out. Once you start working in the industry, you'll start to understand which areas are more interesting for you than others. If all else fails, focus on that last question in the preceding list. If you don't know or don't care one way or the other right now, focus on what brings you joy as a person. People always like to say that markets are saturated. Even before I got started, I was warned of oversaturated job markets by skeptics. But if I had listened to them, I wouldn't have embarked on a journey that showed me how much I was capable of. I wouldn't have had the job experiences I ended up having. I wouldn't have written this book.

I'll tell you what jobs aren't saturated: the ones held by people who are passionate about what they do and find joy in their daily tasks. Do yourself a favor and focus on some area of product or AI that lights you up. Maybe it's related to something you like to do personally – a hobby or a creative pursuit. Let this be reflected early on in your project portfolio or internship experiences if possible. Early on, when I was building my project portfolio, I worked on a data product that compared markets for buying and selling versus renting. At the time, I was deciding whether or not to remain in NYC and I decided to build a product that I could use in my personal life.

I always had an interest in real estate and in making the most of my expenses, particularly one as big as housing, so I decided to let the data speak for itself. I ended up using the tool I built to help me with my own decision-making, which was already great. But then, months later, I was offered my first role in product because of that first project. What started out as an interesting thought exercise led to my first breakout role. At the time, I was focused on remaining in data science. I was actively interviewing for data science jobs and decided, on a whim, to consider a product role with a property tech company that came my way.

My hobby led to me not only building a product I could use in my personal life but it also led to a job opportunity and to the discovery of a whole new job family I've come to love more than any other job I've ever had. So you really don't know what might come of following your personal passions. We all have something we're interested in; something that's well suited for our personality and our talents. It doesn't have to be the be-all and end-all. Real estate isn't even in my top five favorite interests but it was interesting to me at the time. We don't have to get super serious with it; what we do professionally doesn't have to represent the core of who we are. But it helps if the entry point is something that does excite us on some level. It makes everything that comes after that easier.

If you're really not sure where to get started, remember that there are already large concentrations of AI in today's market:

- **Consumer-facing (B2C) AI products:** This is anything that's being developed for end users to interact with directly, like home devices, personal assistants, and recommendation systems. These will be areas where **user experience** (UX) is emphasized. Most of the time, they'll also come with some sort of NLP or computer vision since these will be apps that have to do with your local, immediate environment.
- **Enterprise (B2B) AI solutions:** If you're looking to stay within the confines of B2B solutions, there will be no shortage of work for that too. All businesses need AI solutions to help them with day-to-day operations. Analytics platforms, automation tools, or anything that helps teams with making decisions will be applicable here. If your background is in enterprise or SaaS software, process optimization, or data platforms, you'll find a thriving ecosystem here.

Beyond these two major areas, you also have subgroupings based on industry. We're seeing robust ecosystems for AI PMs in healthcare, finance, transportation, manufacturing, and logistics. Because of the pervasive effects of AI, we're going to see AI applications crop up across all industries. So if you've always wanted to break into any one industry, maybe you can help become one of the innovators to bring that space into the future.

Case study

As I thought about a case study for the chapters in *Part 4*, I was tempted to create hypothetical career trajectories for readers to follow along. But to do that would be intellectually dishonest. It presupposes that there are tried and true pathways to success as an AI PM, and it's exactly what I want to avoid in this book. The truth is that your unique blend of experiences, tastes, and interests will contribute to your success as an AI PM in ways no one can predict. The idea that one can follow a set path or trajectory to land the job one wants is one that I want to leave in the past. When I think of AI PMs today, I think about how much potential there is for AI products and the importance of bringing people into the discipline from various backgrounds.

Here, I will be sharing my own journey to becoming an AI PM and some of the considerations that were relevant to me as I was getting started. The following is a breakdown of the steps I took to demonstrate my abilities as a well-rounded AI PM early on:

1. Learned Python and got a foundation for data concepts through MOOCs (MIT xPRO/Coursera/Udemy/DataCamp)
2. Completed a data science apprenticeship (Thinkful)
3. Built a project portfolio with 7 ML projects
4. Volunteered my data science skills for an early open source app (Fight Pandemics)
5. Wrote articles about my knowledge of data science, ML, and product management
6. Worked as a data scientist for an on-demand gift-giving app (Gesture)
7. Built a mental health-focused NLP app with a peer from Thinkful and ran it through an accelerator with Women in AI (Waterbear/Akeira)
8. Volunteered to organize events locally for Women in Data and Women in AI

9. Started a book club focused on AI and data science concepts for continuous learning
10. An 8+ year career in account management and sales helped me navigate customer empathy and problem-solving skills

Let's take a look at how I was able to demonstrate competency across the top technical and soft skills AI PMs need to be successful:

- **Tech skills**: I was able to demonstrate my technical proficiency by completing online coursework, finishing an apprenticeship, building a portfolio, and writing about my projects online. I was also able to apply my technical skills to Fight Pandemics and Gesture, the two organizations I supported in a data science capacity before my first AI PM role.
- **Business acumen**: For this, I leaned most on my prior 8+ year career in account management and sales for B2B SaaS companies. In those roles, I was client-facing and my primary areas of focus were customer experience and contract retention. In some cases, I had a hard target. This showed a proven client-focused track record of retaining SaaS customers over time. I also built and launched an app with a peer from my Thinkful program, an NLP-based mental health-focused app.
- **Communication and leadership:** I served as the NYC and Boston co-lead for Women in Data, organizing local events that supported the NYC and Boston communities of data and product professionals. I also hosted and moderated a monthly talk with Women in AI where I would find women doing incredible work in AI and talk to them about it, as well as a monthly book club for Women in Data where we would read and talk about a new book about data, AI, ethics, and management. I wrote about all these experiences on my personal blog on LinkedIn and Medium.
- **Problem-solving**: I demonstrated this largely through my portfolio and open source project work. Because each of my projects was devoted to an area of my life I wanted to problem solve, I was able to do this organically as I talked about my projects. Because Fight Pandemics and Gesture were the only relevant work experiences I had before my first product role, I was able to lean on those real-world examples during behavioral questions as well.

Rather than give you archetypes to aspire to, I welcome everyone who's interested in becoming an AI PM to find their own thought leaders and role models and to learn from those individuals directly or indirectly.

Summary

We've covered a lot in this chapter. We started with a discussion of the requirements for AI PMs entering the space from a theoretical and practical perspective, as well as a look into where the AI PM space can be better served by newcomers. We also looked at strategies and tools you can put in place to help you stay supported and inspired. We briefly looked at the major areas of AI that you could get involved in once you are ready to dip your feet in and concluded with a real-life case study to ground these ideas further.

Remember that this is just the first chapter in our discussion on the AI PM career path and we will be building on these ideas further in *Chapters 18* and *19*. We've only scratched the surface of this challenging and rewarding career path that's still very much figuring itself out. AI and tech are here to stay. Even with market fluctuations and periods of uncertainty, we can be sure that our internal operations, ways of doing business, and the overarching systems that govern us will remain digital. We will continue to see data and AI disciplines expand, particularly as we discover new use cases. On some level, building a career as an AI PM is akin to building a career as a futurist. You're tasked with imagining technological futures that have yet to materialize. Dream big and start with what lights you up.

In *Chapter 18*, we will be going deeper into understanding what makes a good AI PM. We will be looking at the various competencies of an AI PM: technical proficiency, business acumen, communication, leadership, and problem-solving. We'll also be exploring some of the challenges and opportunities individuals can face as they're leveling up in this career path.

References

- *Data Science Society*: `https://www.datasciencesociety.net/`
- *Mind the Product*: `https://www.mindtheproduct.com/`
- AI4ALL: `https://ai-4-all.org/`
- *Deep Learning AI Community*: `https://community.deeplearning.ai/`
- *WiMLDS*: `https://wimlds.org/chapters/`
- *Women in Data*: `https://www.womenindata.org/`
- *AI Product Management community*: `https://ai-pm.net/`
- *Women in AI*: `https://www.womeninai.co/`

18

What Does It Mean to Be a Good AI PM?

In the previous chapter, we got our feet wet with what it takes to get started with the AI PM job function. We discussed how people break into the AI PM role, what skills they need to build up to be considered for the role and what ultimately makes them successful when taking on their first AI PM roles. In this chapter, we will be expanding on these concepts and talking more about what builds a strong foundation for longevity in the AI PM career.

We'll spend the bulk of this chapter discussing what it means for an AI PM to be effective. AI PM is a varied job family, and we will be going over the many facets of it to help you best understand what parts of the role you might find most enjoyable, and which parts you might want to work on to develop. We will also be exploring the various ways different people will be challenged by the role based on their previous experiences and how all AI PMs can set up their careers for sustained growth.

Because of the incumbency of the AI PM role, there are more nuances to get into other than technical competence, comfort, and expertise, as we've discussed in previous chapters. If you recall, in the last chapter, we highlighted five areas that AI PMs will have to demonstrate competence in: technical proficiency, business acumen, communication, leadership, and problem-solving skills. We will build on these ideas in this chapter by reinforcing what demonstrating these competencies looks like in a mid-level career. As the chapter goes on, we will cover the following areas:

- The practical stressors and expectations related to the various hats the AI PM needs to wear
- What a proven track record of core competencies looks like
- What communicating accessibly means from the perspective of an AI PM
- Common challenges and opportunities that come as AI PMs rise in their careers
- How an AI PM can take care of themselves as they are expanding

We will finish with a contextualized example of the areas above with our case study example, as well as with a summary of the chapter, as always. The path of an AI PM is challenging but rewarding. Anyone can build a successful, lengthy career as an AI PM, because it's a professional area that will see more growth as we expand with AI proliferation, if they are willing to apply themselves in service to their products, markets, and organizations.

> As we saw in *Chapter 17*, the current landscape of AI PMs tends to fit a profile. But it's my hope that books like this one will inspire you to break into the field even if you don't fit this stereotype and enrich the role with your own unique blend of experiences and values. Anyone can do it if they want it enough, and that means you can too. If reading that gives you pause, because deep down you don't feel you have what it takes, I challenge you to contradict that side of yourself. We all deserve to build careers around what lights us up and inspires us. If working as an AI PM makes you feel excited and inspired, or if it simply sounds like a good way to make a living, that will drive you to find success. For those of you who aren't so sure, the rest of the chapter may very well dissuade you. That's OK and valuable too! It's better to make informed decisions than to never try at all because of imposter syndrome.

A job family of many hats

In this section, we will explore the many practical stressors and expectations of an AI PM based on the competencies we introduced in *Chapter 17*: technical proficiency, business acumen, communication, leadership, and problem solving. I will be using an umbrella for each of these five competencies to keep this section digestible, but you'll notice a long list of hats under each one. We can think of each competency as a concentration of skills and responsibilities and not just as standalone skills. Let's get started!

> Remember that listing these competencies isn't about promoting excellence in each one, but rather is being done to demonstrate the full range of skill sets and competencies you can explore as you grow and mature into your AI PM career.

Technical proficiency

An AI PM's technical proficiency encapsulates various roles that are important when aligning AI capabilities with business goals, ensuring responsible data usage, and maintaining high quality across their product. The AI PM leverages technical insights to solve real-world problems and documents that work to track its impact on the business. Ongoing learning through academic degrees, certifications, and bootcamps can help AI PMs stay updated on evolving technologies and ML models and they can in turn translate complex AI concepts for non-technical audiences and create accessibility through workshops and communications. Technical proficiency isn't just limited to ML models but extends to areas like data management, product quality control, and data analytics, as AI PMs need to make sure the data used in their AI systems, the performance of their products, and usage patterns from their users and customers are all supporting their products.

Technologist

AI PMs aim to build technological solutions to real business problems and can make informed decisions to make those solutions possible. To bridge the gap between technical resources and strategic goals, AI PMs need to have a solution-oriented mindset and a deep understanding of both the product's tech stack and the real needs of the users their product will serve. By applying technical knowledge in ways where AI can genuinely enhance product functionality to improve user experience, AI PMs are able to align their product development toward broader organizational needs and market demands.

Recommended steps: Maintain a personal website, blog, or repository where you can keep track of business problems you've helped solve on a quarterly basis to establish credibility, reflect on your achievements, and refine your approaches over time. Document the business problems you're tackling with AI solutions you're implementing and track the impact of your progress internally as well through user stories and product requirement documents. Share insights with your broader community to build a narrative around your journey as a technologist. Offer evidence of your success in using technology to create meaningful change for users as an AI PM.

AI expert

While you'll be working with data scientists and ML engineers in your work as an AI PM, you'll still need to stay current with advancements, new use cases, and innovations in machine learning models, approaches, and AI methodologies to work effectively with those stakeholders. This helps you make informed, strategic decisions about AI technologies employed in your product. Depending on the business requirements and user needs, you'll be able to demonstrate your understanding of the appropriate models and frameworks that will solve business and user challenges best.

Recommended steps: We covered certifications, bootcamps, and degrees in the last chapter. Get them on a semi-regular basis to keep your knowledge fresh. Structured learning approaches can help you make progress and stay agile in a way that feels sustainable and gradual. Try to set goals for yourself to complete at least one AI-related learning program a year or on a bi-annual basis. Maintaining community and communication with AI professionals regularly through online communities or conferences is also a great way to keep growing as an AI expert at your organization. Sharing or documenting these experiences is also a great way to integrate their wisdom.

Technical translator

One of the essential roles of an AI PM is helping to make technical concepts accessible to non-technical counterparts. Building alignment and clear understanding across teams is essential for enabling stakeholders to grasp the core value and functionality of AI models and product capabilities in plain, accessible language. This helps all teams to foster better decision making and helps make sure everyone from marketing to sales to leadership can meaningfully contribute to the AI product's strategic goals and vision. Writing about the AI-powered solutions in your product will make sure you can refine expectations with all your relevant stakeholders as you're working on AI initiatives on your product journey.

Rcommended steps: Write about the AI solutions you're supporting for a non-technical audience whenever you can, and share those pieces internally where applicable and with your greater network. For teams that are especially struggling, you can think about crafting AI workshop sessions around AI functionality, use cases, and strategic impact so that audiences can better understand AI's role in the product. This can also serve as a space for stakeholders to ask questions and enhance their understanding and confidence in your AI product.

Data steward

In this role, you're responsible for establishing data governance standards and making sure the internal data your models train on is accurate, relevant, clean, and free from bias. Your role impacts the quality and trustworthiness of your product's outputs, as well as the model's accuracy, fairness, and compliance with regulatory standards. The integrity of your training data allows you to mitigate risks related to biased, incomplete, or inaccurate datasets.

Recommended steps: Write about your governance standards and the challenges you face making sure the internal data your models train on is accurate, relevant, clean, and free from bias. Communicating regularly to your leadership team about specific examples where data management significantly contributed positively or negatively to real business outcomes and user experience strengthens the case for data stewardship as a vital component of the AI lifecycle and reinforces the importance of investing in proper data management.

Data strategist

While a data steward might focus on the day-to-day challenges in the data, the data strategist role encourages a high-level approach to data management by developing long-term strategies for the collection, management, use, and deployment of data in production environments. Building strategies around effective data practices not only strengthens the performance of your product's AI capabilities and features but also optimizes the data pipeline itself. Data strategists concentrate on the strategic, long-term planning of data acquisition and use.

Recommended steps: Organize visible data strategy sessions regularly and ensure key stakeholders and leaders are involved from the start. Maintain documentation for data strategies that relate to the collection, preprocessing, transformation, and QA for your organization's data. Make sure it's accessible for all relevant stakeholders to understand data handling protocols and lineage in a centralized place where they can ask questions and bring up issues to establish a broader culture of data fluency, particularly as business objectives evolve.

Quality controller

You're the last line when it comes to making sure your product meets quality and performance standards before it's shipped. An AI PM makes sure that all AI features and functionality meet the high standard for quality and security before it's launched. This means being fluent in rigorous testing practices like stress, load, and usability tests that come along with continuous quality assurance. The AI PM must make sure that the product is ready for launch by assessing it for technical standards and user expectations, particularly for AI products where data and user inputs can vary significantly.

Recommended steps: Share continuous product improvement strategies and QA best practices with key stakeholders and leaders involved. Keeping stakeholders informed of testing results and QA can show how improvements contribute to the product's value. Include edge case testing, bias analysis, and fairness monitoring in your QA best practices.

Analyst

As an AI PM, you're regularly using product analytics and profiles of user/customer behavior to make informed decisions about what's truly resonating with them. You're consistently monitoring and interpreting changes in usage patterns, engagement metrics, and feedback trends to provide actionable insights into the features or functionality of your product that makes the biggest impact on the users you're serving. This helps you adapt your product to the evolving needs of your users and market.

Recommended steps: Regularly update your organization on trends, KPIs, and metrics that are being used to gauge the success of your AI product. Share reports on these areas with cross-functional teams and leadership to showcase user engagement patterns and areas of improvement to ensure you have buy-in for course corrections later on. Use analytics to measure how well AI features are doing and how they're impacting the product's strategic goals and vision.

Business acumen

This aspect of the AI PM role focuses on the functions that connect the product to the business landscape and market it's serving. An AI PM is responsible for defining the product vision and aligning it with the company's goals by emphasizing the product's financial impact and market performance. A big part of guaranteeing that performance is building partnerships and collaborations around their product to ensure it has the proper distribution channels for success. The business side of the AI PM role also requires a deep understanding of the market conditions and readiness for the product they want to build or maintain. This means AI PMs need to be well versed in the competitive landscape they're playing in to understand how to best position their products and inform future innovations.

Strategist

In this role, an AI PM is defines the product vision and strategy and makes sure it's aligned with the company's overarching goals. This is a pivotal role because it shapes the direction the product is going in and articulates that direction for stakeholders to understand. Articulating a clear, compelling vision and creating actionable strategies that inform product development creates direction for the team and inspires stakeholders across the organization to believe in the product's unique value proposition and how it addresses the needs of the market and its users and customers.

Recommended steps: Set up quarterly or semi-annual town halls where you publicly deliver the company's product vision and strategy and explicitly connect those to the company's overarching goals. Maintain a regular biweekly or monthly product strategy meeting with leadership or your GTM team to review progress, discuss market trends, and adjust strategies when applicable. Build market analysis and feedback into your strategy discussions to analyze competitors, assess insights from users, and understand industry shifts.

Revenue driver

In this role, you're not only making sure the AI product is having an impact on the business's bottom line but also that the product is meeting its revenue targets. By building a narrative around financial success and aligning that with KPIs and metrics that demonstrate the revenue the company is expecting, you're demonstrating your effectiveness as an AI PM to your business.

Recommended steps: Make sure you're meeting regularly with leadership or GTM teams to demonstrate revenue growth to show your AI product is having a positive impact on the business. Leverage data analytics to track sales growth, revenue per user, sales conversion rates, customer retention, customer acquisition costs, and customer lifetime value.

Partnership builder

In this role, you're establishing necessary partnerships that will positively impact your product, mitigate risks, and contribute to the commercial success of the business. Identifying and cultivating relationships with external partners like technology providers, research institutions, and industry leaders creates synergies that can benefit your product's development and market reach. Effective partnerships provide access to additional resources and expertise your organization otherwise wouldn't have and can lead to new innovations on product features and functionality that make your product more competitive in your market.

Recommended steps: Identify strategic opportunities and potential partners whose capabilities are complementary to your product. Showcase strategic partnerships that have enhanced your AI product offering through blogs, case studies, or white papers. Documenting successful collaborations strengthens partnerships and validates the benefit of collaboration.

Innovator

You're always looking to identify new ways AI can improve your product and reinforce a culture of innovation within your product organization. Here, your primary focus is on continuously exploring and finding new ways that AI can enhance your product's capabilities and market positioning. It also means finding opportunities for resources and curiosities to be expanded within your organization with the resources and team you already have.

Recommended steps: Celebrate the top recent innovations for the month, quarter, or year through blogs, articles, or public talks. Bring these innovations up as examples during strategy meetings with leadership and highlight initiatives that align the product vision with the company's strategic goals. Look for ways team members can be empowered to experiment by providing resources, time, and support for experimentation in the form of hackathons or innovation sprints.

Market researcher

AI PMs need to understand their market to tailor their product offerings in a way that meets customer demands and expectations. In this role, you're consistently doing market research to understand the relationship between AI innovations, current offerings, and user/customer pain points so that you can develop a product that's relevant and competitive.

Recommended steps: Document your market research activity through reports you can share internally and externally to establish yourself as an expert in your domain. Include successful market positioning strategies and outcomes that have worked positively for your product. Directly engage with your target audience where applicable through online platforms and communities or through direct user feedback and product surveys.

Competitor analyst

While a market researcher analyzes various components of the market, here you're focusing on the competition by analyzing your competitors' strategies to maintain awareness about what makes your product competitive in order to understand how you can innovate beyond them. Continually assessing the strengths and weaknesses of competing products allows you to maintain a competitive edge in a rapidly changing AI landscape. This analysis doesn't just allow you to refine your product offerings but it equips you to anticipate market shifts and respond proactively.

Recommended steps: Document your competitive analysis and highlight areas where your product is strong compared to your competition and areas where it's weak. Share it with your leadership team whenever there's a significant update to it. Conduct a SWOT (strengths, weaknesses, opportunities, threats) analysis for each key competitor to understand how your product can remain competitive. Establish performance benchmarks by comparing your product's key metrics against industry standards and competitive offerings.

Communication

Effective communication is important in most functions, so for the AI PM, this facet of the role is crucial. AI PMs have to collaborate with a number of cross-functional teams to keep their projects, roadmaps, and sprints on track. Communicating all major AI features and capabilities throughout the product lifecycle and fostering transparency means stakeholders are involved and considered in a meaningful way. All significant product changes and updates have to be communicated through the AI PM, and when those updates are not fully understood or appreciated, it's up to them to step in and address the gap in whatever way makes the most sense. Communication doesn't have to just be functional; it also serves a strategic purpose for the AI PM. Communicating risk and mitigating potential losses is essential to keeping lines of communication open between PM and leadership teams.

Project manager

AI PMs need to have strong organizational skills in order to coordinate with cross-functional teams to make sure projects, roadmaps, and sprints stay on track. Bringing in perspectives from engineering, design, marketing, and sales and integrating those perspectives through regular check-ins, status updates, and sprint planning meetings keep alignment strong across product activities. The AI PM not only makes sure alignment is kept across various product areas but that these influences are managed in a timely manner that adheres to the roadmap and schedule that has previously been laid out. The focus here is on managing resources in a way that tasks stay on schedule. In many cases, AI PMs also conduct post-mortems after a significant sprint or epic has been completed to enhance their work in the future.

Recommended steps: Make sure you're conducting post-mortems for all significant projects to celebrate successes and establish lessons learned for next time. Document all major AI features/capabilities/products as you go. Use collaborative tools like Jira, Trello, or Asana to track progress toward your sprints/epics and share updates.

Change agent

Ultimately, you're responsible for outgoing communications within the company and outside of the company when it comes to product updates and pivots. Here, your role is not only to drive the development of AI features and capabilities but also to keep your stakeholders engaged throughout the process. You represent a leader in your organization and your influence has the ability to impact meaningful change on how stakeholder teams work together and recognize one another's accomplishments.

Recommended steps: Regularly and publicly recognize when others contribute to product success, particularly after a major product pivot or update. Encourage feedback from stakeholder teams by regularly seeking it after major releases or initiatives.

Stakeholder manager

As an AI PM, you're the point of escalation from any stakeholder team your product organization works with. You're the main point of contact for product changes and updates to all stakeholders; you set and manage their expectations on an ongoing basis. When any of the major teams have an issue or concern they need to raise with regard to your product, they will likely need to raise it with you directly, which means you will have to balance the needs and priorities of all the key players within your organization. This role is long-term focused in nature because many of your stakeholders will function like partners to your product organization.

Recommended steps: Communicate how you're setting and managing stakeholder expectations during 1:1s or meetings with leadership. Share a detailed product roadmap that clearly highlights your product's vision, strategy, and significant milestones. Keeping stakeholders well informed and proactively engaging with them about important developments will help you align their interests with the product vision and foster a supportive, collaborative environment. Nurture long-term bonds with collaborators and stakeholders in ways that are authentic and supportive.

Educator

AI PMs can help manage, facilitate, or plan product training and support for team members and stakeholders that have gaps in communication. Bridging knowledge gaps in your product's features, functionalities, and best practices for team members and stakeholders ensures that they have the proper context when they're contributing to your product, no matter what their role.

Recommended steps: Organize/conduct product training given to stakeholders. Make sure that training is documented and shared with others at the organization who couldn't join. Create a public glossary or index of terms/features and how to speak about them to customers/users. Create, document, and share training materials whenever applicable.

Risk assessor

AI PMs need to identify, communicate, and mitigate all risks associated with their product to their stakeholders and leadership team. By keeping stakeholders and leadership aware of potential challenges and threats to your product, you're not keeping them in the loop, you're demonstrating your stewardship and responsibility as a PM by helping them ease the burden of leadership. Without having all the facts, leaders can't make informed decisions, so this aspect of the AI PM role focuses on the importance of delivering critical news and solutions to address it.

Recommended steps: Include a risk assessment section for your product strategy sessions with leadership and GTM teams. Have a quarterly or ad hoc meeting with legal and share findings with your leadership/higher-ups. Document a risk mitigation plan for major AI features/capabilities/products as you go. AI PMs need to develop strategies to minimize risk for their products so when these do come up, assess how to prevent them in the future as well.

Leadership

One of the critical functions of an AI PM is the ability to lead. AI PMs have to translate their product vision internally and externally through various channels. They also have to foster collaboration and manage various facets of a team whether they directly manage people or not. Often, they're sharing knowledge, communicating the story of the product, and integrating their product vision, strategy, and value proposition continuously in different ways. AI PMs also have to keep various stakeholders and teammates motivated toward an overarching goal or purpose and regularly need to demonstrate competency around ethical AI considerations. A leader is characterized by their ability to guide and inspire others toward a common vision and the AI PM role is not short of opportunities for this.

Visionary

You translate your product vision within your organization, both hierarchically and horizontally, as well as outside of your organization through product marketing, thought pieces, white papers, blog posts, and public talks. This involves not only defining what the product aims to achieve but also ensuring that every team member, stakeholder, and external audience understands and aligns with this vision.

Recommended steps: Maintain a regular, monthly cadence of product marketing, thought pieces, white papers, blog posts, or public talks you're giving related to your product and its ability to solve your customers' biggest issues.

Ethicist

AI PMs have a responsibility to ensure their AI systems are developed and deployed in ways that are ethical and up to standard. This means you ensure you mitigate issues of bias, fairness, or transparency when they arise, and that you're you keep community of users and customers aware of these standards regularly as a part of your product marketing efforts. Fostering an ethical framework reflects a commitment to social responsibility and it's a part of the brand identity around your product. Ethical concerns are of increasing importance to consumers, users, and customers, and AI PMs will need to become more vocal about these considerations as time goes on.

Recommended steps: Demonstrate thought leadership by writing about how your AI product is being developed and deployed ethically and how you mitigate issues of bias, fairness, or transparency when they arise. Share these pieces with your network and peers. Encourage broader discussions on the importance of ethics in AI and engage in these conversations to position yourself as a knowledgeable advocate.

Team leader

You act as a team leader by making sure all stakeholders are aware of their role and their importance to the product team. This involves creating an inclusive environment where team members feel empowered to approach you with questions or concerns regardless of your direct management role. Keeping open lines of communication fosters trust and encourages people to collaborate in a productive, dynamic team.

Recommended steps: Communicate to your leadership team or direct manager about how you're leading teams, fostering collaboration, or resolving conflicts when they arise. Highlight how you facilitate teamwork and share updates on team performance, challenges faced, and goals achieved.

Storyteller

As a storyteller, your role is to regularly communicate the product vision, strategy, value proposition, and the benefits those things will bring to stakeholder teams. You also tell the story of your product and what it can do for your market, customers, and users. Good AI PMs know that effective storytelling is about clearly articulating narratives that resonate emotionally with audiences to keep them invested. By connecting the benefits of your product in relatable and compelling ways, storytelling can inspire action for internal teams and users alike. As you weave your stories into various pieces of product collateral or promotional opportunities, you'll refine how naturally and authentically you deliver your message.

Recommended steps: Regularly communicate the product vision, strategy, value proposition, and the benefits those things will bring to stakeholder teams with every major sprint, product update, or pivot. Enhance your storytelling skills by weaving your storylines into various communication channels like presentations, newsletters, customer-facing changelogs, blog posts, meetings, formal talks, podcasts, articles, happy hour conversations, or anything else applicable.

Motivator

When the team loses inspiration, when they forget what they're doing and why they're doing it, you're the one who's going to help them remember what their hard work is all for. Drudgery and overwhelm can set in when you're working a full-time job. Without the influence of someone who can bring the team back to the core purpose of their work, stagnancy and disillusionment can set in. By keeping the big picture alive, you remind everyone of the impact and meaning behind their efforts. Encouragement and perspective can reignite drive and help others overcome obstacles with renewed purpose and optimism. This is one of the more emotional aspects of the job and good AI PMs have the emotional intelligence to recognize when this important ingredient is missing from the team.

Recommended steps: Affirm the product vision and goals when others seem lost, frustrated, disillusioned, or confused. Acknowledge individual and team contributions publicly or privately, as appropriate, to strengthen morale and reinforce the value of their hard work.

Knowledge sharer

You'll lead by example by creating a culture of continuous learning and knowledge sharing regarding AI, the market, or areas related to your product. Because successful working AI PMs are already in the process of acquiring and reinforcing knowledge for their own personal excellence, this role is about sharing and applying that knowledge with others. The goal here is to disseminate information. By actively finding moments to share insights, resources, and product knowledge, you're making sure your team and stakeholders remain well informed and empowered to make their own decisions. This contributes to overall growth and expertise within your organization and it fosters a lot of trust and belief from your coworkers.

Recommended steps: Maintain a public calendar of knowledge sessions, training, and resources, and make public announcements when there has been a contribution to your product's knowledge base.

Problem solving

This aspect of the AI PM role encompasses the various responsibilities of the role. Solving a problem is about considering a variety of potential solutions and bearing the responsibility of settling on one choice that will bring the most benefit. AI PMs have to sift through a variety of problems ranging from customer feedback to regulation and compliance; necessary technical adaptations to general conflicts that can arise. There's no shortage of problems at a company, but there is a shortage of effective problem-solvers. Many organizations face challenges related to communication breakdowns, resource limitations, and evolving market demands. Building your problem-solving muscles to navigate complexities, identify root causes, and enable collaboration will serve you in many situations.

Customer advocate

As a customer advocate, you're often placed in a position where you're sifting through various channels of customer feedback and looking for opportunities to optimize product development in service to the major trends your users and customers are consistently delivering back to your product organization. You prioritize user needs and pain points to ensure your product has a real impact on what's most meaningful to them – to make sure it's genuinely addressing enhancements that would positively impact their lives. You empathize with the end user's journey and craft an AI product experience that meets their needs and fosters loyalty through consistent, value-driven improvements.

Recommended steps: Develop compelling case studies that demonstrate how your product delivers value to your customers/users. Collect and present major user feedback, reviews, testimonials, and survey findings to stakeholders to reinforce product alignment with user priorities and build evidence-based support for future product iterations.

Regulatory complier

While your organization might have legal teams that are more heavily concerned with regulatory oversight and adherence to frameworks for auditing purposes, you'll likely need to work with those teams as you're increasingly integrating AI into your product experience. This means staying familiar and up to date on legal standards and regulatory frameworks in order to ensure your work aligns with them and doesn't pose potential risks. By staying informed about relevant and timely laws and regulations that relate to your product, you can also identify opportunities that can add value to your product and give it a competitive edge.

Recommended steps: Regularly stay informed about relevant and timely laws and regulations that impact your product; be the first to flag potential risks and opportunities that compliance with those regulations poses to your product's success. Document insights and potential regulatory or legislative milestones you're anticipating in the future that could impact your product decisions or strategic direction long term.

Facilitator

We've discussed *knowledge sharer*, *stakeholder manager,* and *project manager* previously in this list, but none of these gets to the heart of what a facilitator really does. A facilitator's primary focus is to create an environment where collaboration can thrive across diverse teams and to remove any barriers there might be to that end goal. You regularly align with stakeholder teams to prevent friction as much as possible and problem solve as to how to prevent it from happening in the first place. Facilitators guide discussions to make sure all perspectives are heard and that there are no gaps in necessary conversation streams when teams do collaborate and meet. Facilitators also help the team work through misunderstandings before they become serious threats or blocks. Proactively managing group dynamics, identifying areas of misalignment, and encouraging open dialogue may be aspects of other roles, but they're not the primary focus.

Recommended steps: Maintain a regular cadence with all major stakeholder teams to regularly keep them up to date on product goals, milestones, and changes. Do not meet with them ad hoc. When meetings do occur, make sure the agenda is representative of the goals of that meeting and that all relevant voices have the freedom and safety to speak openly.

Data-driven decision maker

Earlier, we covered the *analyst* role, which focuses on tactically finding trends in data concentrated on user behavior and engagement metrics. While this is a skill set in and of itself, acting on those trends and insights and weighing options for how to move forward is what this role is all about. You use data and metrics to make the right decisions for your product and its positioning based on competing priorities from various stakeholders. Aligning decisions with overall product goals and priorities means the product will evolve based on comprehensive data insights. In this role, AI PMs facilitate discussions with leadership and stakeholders about decisions that will shape the direction and positioning of the product long term.

Recommended steps: Maintain a dashboard or deck that is regularly updated with new data points that relate to product KPIs and metrics. Use this deck when you're presenting to stakeholders and leadership. Track the most relevant KPIs, metrics, and data points that reflect product health, customer satisfaction, and user engagement to justify key decisions in terms of measurable impact.

Adaptability manager

AI PMs have to be adaptable and flexible in the face of changing conditions, whether they come from internal resource constraints, the market, or technological advancements. You remain adaptable and ready to pivot whenever the product needs to take a different approach. As the product representative who has to regularly engage with stakeholders to convey adaptations, this means you'll be the one to provide context and explain why certain changes were chosen over others. In this role, you're not just deciding on which pivots to make. You're also defending the choice itself based on your vantage point, which balances leadership's vision with capacity shifts.

Recommended steps: Clear all adaptations and pivots with your leadership/GTM team when you regularly meet with them. Then, communicate those agreed-upon adaptations to all your key stakeholders when you regularly meet with them. Document product adaptations and changes to key goals and strategic initiatives; keep track of the rationale behind those choices so that you can reflect on them for future decision making. Regularly review industry reports, articles, and analytics to remain proactive to external trends.

Conflict resolver

While the *facilitator* role focuses on aligning teams and ensuring smooth frictionless collaboration, the *conflict resolver* is focused on addressing and mitigating disagreements and conflicts related to the product. Conflict can manifest in various forms. Whether conflict is immediate and arises temporally during meetings or arises from a team's ongoing tensions and competing priorities, you're the first point of contact when disagreements or issues come up. Often, the two intersect. A temporal conflict in a meeting could reveal a deeper issue related to competing goals that have not previously been addressed. Differing opinions, misunderstandings, or misaligned objectives need to be weighed and considered with a problem-solving, solutions-oriented approach. An AI PM with strong conflict resolution skills knows when to step in to address deep-rooted issues to prevent more frequent conflicts later on.

Recommended steps: Keep track of major disagreements or prioritization issues that come up related to your product; flag those you think are strategically relevant to your leadership/GTM team during your regular product strategy sessions. Learn about conflict resolution techniques on a yearly or biannual basis to empower conflict resolution attempts within your team and with other stakeholders. Build practices and protocols around conflict resolution that allow team members to know how to escalate issues that they feel are unresolved or too significant to handle at their level.

That's a lot of hats. There are likely many more that are missing from this list. Perhaps we can make this list longer for the third edition, at the risk of scaring off more curious, prospective AI PMs out there. Though this list is quite long, it adequately visualizes both the impact and the complexity of the AI PM role. It also highlights the important point we made in the last chapter: technical expertise does not make up the bulk of the role. I would consider it more of a prerequisite. Most of the items on this list are non-technical in nature. At its heart, the AI PM role is about critical thinking and people. Next, we will explore how we can use critical thinking and interpersonal relationships to excel in the AI PM role.

One of the reasons AI PMs struggle the most within their roles is that others are often not aware of what they are doing in service to their products. We'd all like to live in a world where we're rewarded and celebrated for the work we do without having to take on the burden of making it apparent to others.

But we don't live in that world. If you're not able to visibly position the work you do, the day-to-day stressors and challenges of your higher-ups, leadership team, stakeholders, and professional network will blind them to your effectiveness. We try not to take this personally. But if we don't understand this point, we may feel that we're doing quite a lot without having the necessary support and recognition that will fuel us for the long haul. Navigating the core competencies and championing your contributions to the business through them will contribute most to your real and perceived success as an AI PM.

These actions might seem overwhelming at first. But this reinforces the importance of structures in the day-to-day life of an AI PM. Organize activities that will make the biggest contribution to your real and perceived effectiveness as an AI PM into categories that are weekly, monthly, quarterly, and yearly. I've done my best to include how often you should be checking in with various teams and stakeholders in the bullets above, but ultimately the decision is yours. You know your leadership, GTM, and stakeholder teams best. Just remember it's better to book time and decide you don't need it later than to struggle to find time with people. Be proactive with the structures you set up in your AI PM role. Let these structures work for you. One of the major structures you'll set up to enable success in your role is communication channels. We'll explore what these look like next.

The AI whisperer and the role of communicating accessibly

I can't stress the role of making AI accessible to the various non-technical stakeholders you will work with enough, so this deserves its own section. As an AI PM, much of the product communication will fall directly on you. You may have large product teams at your organization, but that doesn't mean this role won't fall on you. Whether you work for a large or small company, despite the size of the product team you're working in, if you are responsible for communicating your product vision, strategy, and roadmap to your organization, you will bear this burden.

The reason for this is technical teams are often communicating, troubleshooting, and problem solving internally. It's rare that major stakeholders will go to engineering, data science, or machine learning teams directly to understand what's happening technically with the product. Often, there are no other interdisciplinary or intermediary roles that exclusively deal with making technical concepts accessible. Even if you did have engineering hours to throw at this, they may not know how to explain what they're working on in a language that non-technical counterparts can understand.

This is why the responsibility falls on the AI PM. You serve as the communication nexus that bridges the tech side with the business side. The impact of this is huge because you're the mouthpiece that speaks to leadership, stakeholder teams, customers/users, prospective customers/users, competitors, peers, and curious strangers on the internet. You're also the empathetic, compassionate ear that takes in the criticisms, confusions, delights, and pain points all these various groups have about your product and its perceived utility/value in the market.

Because so much of the product role has to do with communicating and listening, the most advantageous thing you can do in your role as an AI PM is to learn how to speak and how to listen. Internally, this will be an intuitive exercise as you understand what kind of terminology to use as you educate, evangelize, and explain your product's value to others. Because the AI PM role is a nascent role, the responsibility to not just communicate your product's effectiveness but the general effectiveness of AI falls on you. Not being able to do this well will have a significant impact on your product's success.

Imagine your marketing team working on product marketing collateral that confuses what part of your product is using AI and how it works. What if your sales team can't navigate questions or concerns from prospective customers? Or if your customer success reps couldn't articulate how a certain feature works to your users? Worse yet, what if your CEO gave a public talk on your product at a conference and over-amplified product capabilities? I'm sure no CEO has ever done that, but it's a possibility we have to prepare for nonetheless.

Because the AI PM role is so collaborative and cross-functional in nature, hedging yourself against misinformation, confusion, disinformation, and resistance plays a big role in your success. It's a role that will require a lot from you because you have to live and breathe your product in order to meet its demands. Building a career as an AI PM, particularly after the first few roles, is about demonstrating a sustained capacity and appetite for product success. You won't be trusted and supported as an AI PM if you can't mature your capacity for hearing your audience and meeting them where they are.

Common challenges and opportunities as you're leveling up in your career

Having a clear plan of what you want out of your AI PM career and where you want to go is important for anyone looking to build a long, sustained career. You will be inspired to go in different directions as you continue down the AI PM path, but having a clear path (or roadmap if you will) for your career direction will help you harness opportunities and prepare for challenges you'll face as you grow your career. As we've discussed in previous chapters of this book, the following are some challenges and opportunities you may face as you scale your career so that you can best anticipate them as you grow your AI PM career:

- **Tech landscape:** One of the greatest hurdles with the AI PM career path is staying up to date on the fast pace of advancements in AI technologies. It's demanding and many people face imposter syndrome over it because they might have some familiarity only with certain methodologies or models. Then, because they're an "AI PM," they may feel that they need to be aware of all areas of AI to embrace the title with confidence. But this isn't the case. Yes, you'll want to stay up to date with what's going on as much as you can, but you're also working within one area of AI, and mastering that will likely be enough for your product and company. It's often hard for people to balance the need for deep technical knowledge with a general awareness of other AI models and capabilities. It can be solved by knowing what area of AI you want to specialize in and going deeper with that.
- **Alignment:** Maintaining alignment with stakeholders is hard without knowing what each group's intentions for the product and level of AI awareness are. This poses a hurdle for a lot of AI PMs because it's what most contributes to breakdowns in communication. Build trust with these teams first so that you can understand what their goals are and what their understanding of the product's value is. Once you have that, you can use that feedback to inform the roadmap you'll present to the company. If you take every step of this as a standalone step, it will prevent miscommunications later. You can't set realistic expectations with others without first assessing their baseline expectations. Because you're an AI PM, you're assessing their understanding and expectations of both the product you support as well as what AI can do. Both need to be addressed and clarified.
- **Market:** The AI landscape is competitive and will continue to be so as more and more players enter the AI game. This is also true of PMs coming into the space with increasingly more familiarity with AI. You don't exist in a void – you have people around you who should be incentivized to help you. Try to combat the pressure from the market by digesting the findings you get from your research and then build strategy sessions and workshops around them. Your work as an AI PM is to bring all the relevant voices into the room and to receive guidance from your leadership team. They're the ones who took on the responsibility and investment to build a company and play in the market. Let them lead you.

- **Resources:** We've covered the costs of maintaining an AI system earlier in this book and often, as an AI PM, you'll make certain calls based on budget or time constraints. Setting out to achieve unrealistic, ambitious AI product goals is a good way to evaporate your credibility as an AI PM. Just as you build trust and gather AI awareness with the stakeholders above, you very much set out to do the same with your data science, machine learning, and development teams. Try to establish a baseline of awareness of technical teams' capacity early on to understand what velocity they have. Avoid making any commitments until you've gotten to know what your technical teams are capable of.
- **Ethics:** Many AI PMs may feel that assessing their products for bias, drift, and fairness fall on data science teams to flag up but that's irresponsible. Yes, your data science and machine learning teams should feel safe to voice concerns they have around training data bias or issues with model performance across a variety of user demographics. But you can't leave that up to them. You're supposed to be the expert on your product and if you don't set up a plan for catching potential issues that relate to bias and ethics proactively, it could get lost in the cracks. Also, transparency is becoming increasingly important to customers and users. Without running a full assessment of how your product's performance changes based on a variety of factors, you won't be able to give a transparent account of what makes your product effective.

We've covered a lot of areas where AI PMs can expand as they're maturing in their careers over time. For those starting out, it can seem overwhelming to have to think about so many considerations for one job. And it is! We won't all be excellent at every aspect of our role all the time. Depending on the company you work for or the expectations for that particular product role, certain aspects that we've covered so far will be stressed and appreciated more than others. As you mature as an AI PM, you'll build strategies that will help you find the right mix of company, role expectations, and personal preference that will make you shine the brightest. Growing our careers sustainably, as joyfully and gracefully as we can, is the most important job of all. In the next section, we will discuss the biggest contributor to sustainable career growth: self-care.

The importance of self-care

Having a support system of self-care is crucial when dealing with the wear and tear of a career in AI PM (and any PM role really). As we've seen with the first two sections of this chapter, there are many hats to wear in this line of work and a long list of recommended actions to take to meet the needs of these various roles. Without the proper infrastructure to maintain capacity, productivity, and inspiration, people can burn out very quickly because it's often a role with high visibility, expectations, and stress.

This is largely because the investment in an AI PM is high. A lot is at stake for the business and they need to find someone that's truly up to the task. Remember that as an AI PM, you're responsible for the organization's ingestion and understanding of AI, the AI program itself, the product strategy and vision, the execution of that strategy, and, of course, user adoption and satisfaction. These would be demanding projects even if one person was in charge of each, let alone one person being responsible for all.

Taking measures to not only keep your focus and energy levels but also to maintain your emotional engagement will keep you replenished. Often, it's not the actual demands of the job but the emotional resistance that makes us more tired. Combat this by taking on roles where you work closely with direct managers, direct reports, and leadership teams that are supportive and inspiring when you can. I realize this isn't always possible, but how well you work with others will have a direct impact on your success as an AI PM. The more emotionally supported you are, the more mentally clear and capable you will be.

Working with people who disempower or undermine you will have you working against a strong current of resistance. Priding yourself on logic and reason will only take you so far. Eventually, the demands of the job will outweigh your unemotional veneer and you'll find yourself struggling to find the motivation to keep your job. Do yourself a favor early and build supportive teams around you to help you maintain a strong emotional sense of belonging and purpose. Our commitments are emboldened when our minds and hearts are in it for the right reasons.

Maintaining strong mental and physical health is important for all people, but for the AI PM role it should be looked at as a prerequisite. The role of an AI PM is high pressure and demanding, so make sure you're able to manage stress effectively to prevent burnout. There's an established connection between mental health and physical health. You won't be doing yourself or your product organization any favors by sitting behind a screen all day, going from meeting to meeting back to back, and never seeing the sun. You can stay ahead of your health by focusing on the following:

- Incorporate regular exercise like walking, yoga, or going to a gym/workout class to relieve stress and stay fit.
- Balance your emotions and your focus by eating a balanced diet regularly throughout the day.
- Get enough sleep! Don't work late. Don't work early. If you're managing your schedule properly, a 40-hour workweek is enough to do your job effectively.
- Meditate! Create distance around your to-do list and revisit it from a bird's-eye view every day.
- Take mental breaks throughout the day. See the sun. Don't schedule back-to-back meetings.
- Build strong emotional ties around you. See your friends, family, or therapist regularly so that you have a way to blow off steam. No career is worth the dis-ease that comes from isolation and overwork.
- Maintain hobbies, personal projects, and third-space involvements. We are not just our jobs; we're so much more than that.
- Prioritize your work so that you're focusing on what's most important or time sensitive first.
- Delegate what you can.
- Don't work when you're on vacation.
- Define your working hours and stick to them.
- Make professional connections virtually or in person.
- Find an AI PM mentor, someone you trust that you feel has your best interests at heart.

Self-care is self-love and it's hard to build credibility and authority if you're worn thin. Remember that this is advice for sustaining a career, not a job. You can work yourself to the bone, scheduling every meeting you can and staying late often for one job. Maybe two. But eventually, something's going to give. And that something is likely going to come at the cost of your personal relationships or sanity. Give yourself the gift of time and space as you're building your career and learn from the areas that have stripped your capability and focus in the past.

Creating a product vision and strategy, executing that strategy, and communicating both high-level and detailed aspects of your product to every major stakeholder is a lot. These are vastly different skill sets that are expected from one person, so give yourself a lot of grace with this career. It's very understandable if you find yourself being overwhelmed by the magnitude and complexity of the AI PM role. The kind of people that often gravitate to it are the people who have an incessant need to be constantly challenged. But that drive and grit has its limits. Know your limits and protect your mind, body, and heart at all costs.

Doing so will mean that you can start any AI PM role strong. The better boundaries you have, the stronger your capacity can be for managing the day-to-day effectively. Even with starting strong and maintaining your capacity, the reality is that you'll eventually hit turbulence. Things may go horribly wrong. A release might have come with unintended bugs that only manifest themselves in production. Your entire platform could crash unexpectedly. Sometimes, bad things happen even when we do everything we can to prevent them; that's life! Taking care of yourself and putting strong boundaries in place is about building resilience so that you're strong enough to handle the challenges that will eventually come.

Case study

Drawing from my own professional experiences, here is a list of actions I've taken at various points in my career so far:

- **Technical proficiency**:
    - **Technologist:** Maintain articles on my LinkedIn and medium pages and take on speaking engagements to communicate my work. I also moderate events with Women in AI, giving a platform to other women doing incredible work in AI (and learning from them).
    - **AI expert:** Got various certifications and completed two boot camps. I buy Udemy courses occasionally to practice with different ML models outside of work.
    - **Technical translator:** Organized training with stakeholder teams to get them up to speed on mature data practices and I've done workshops explaining models used/alluded to when doing competitive analysis on other products.
    - **Data steward:** Documented data hygiene, data lineage, and data cataloging tools; I've also created data governance and access standards.
    - **Data strategist:** Established documentation for data strategies that relate to the collection, preprocessing, transformation, and QA for my organization's data. Made sure it was accessible for all relevant stakeholders.

  - **Quality controller:** Analyzed top QA issues for various product rollouts and a plan to mitigate them for future releases.
  - **Analyst:** Organized a recurring product strategy meeting with leadership and top stakeholders sharing KPIs and metrics that positively impact revenue and user retention.

- **Business acumen:**
  - **Strategist:** Maintained a regular biweekly or monthly product strategy meeting with leadership or the GTM team to anticipate future product pivots.
  - **Revenue driver:** Gave a weekly product update during a Monday AM all-hands recurring call on how a product was impacting revenue numbers.
  - **Partnership builder:** Won AWS Marketplace partnership that positioned our product as an AWS strategic partner, giving our product higher visibility in our market.
  - **Innovator:** Included a "top 5" in every quarterly product review meeting highlighting the top features/capabilities that positively impacted revenue growth or user satisfaction.
  - **Market researcher/competitor analyst:** Created a market research document that I shared publicly with all major stakeholders that included the state of the market, the major players, our top competitors, and a path for major milestones we were headed for based on our product vision and strategy. I updated it quarterly.

- **Communication**:
  - **Project manager:** Ran sprint planning meetings and set up a process where I was flagged anytime there was a delay with story points for each engineer.
  - **Change agent:** Set up a Slack channel called "Shoutouts" for publicly recognizing the actions/contributions of others and used it to reinforce positive behaviors or pivots.
  - **Stakeholder manager:** Kept a weekly updates meeting with all my major stakeholder teams and referenced the product roadmap for each. The meeting covered challenges from the last sprint and upcoming changes/improvements from the next sprint.
  - **Educator:** Created a data stewards program where data-minded people from each stakeholder team would train on data and AI best practices they could share with their teams.
  - **Risk assessor:** Created a case study assessment for the downstream effects of our platform, which ended up being a white paper of findings we shared with prospective customers that were risk averse.

- **Leadership**:
  - **Visionary:** Wrote regularly in our company blog under the Product tag and represented my work at industry functions and conferences.

    - **Ethicist:** Held workshops for customer success and sales teams about the impact of weights and hyper-parameters of different ML models and how that impacted the customer/user experience for various verticals we supported.
    - **Team leader:** Regularly asked for advice from my direct manager about how I was managing team responsibilities during my 1:1s.
    - **Storyteller:** Formed a data leadership committee and enlisted the support of stakeholder team leaders to evangelize new cultural data maturity practices.
    - **Motivator:** Became the trusted person peers and teammates came to when they were overwhelmed and frustrated and never used it against them.
    - **Knowledge sharer:** Established a feedback loop where anyone could flag an area they were unsure about for clarification; some interesting workshops came as a result.
- **Problem solving:**
    - **Customer advocate:** Regularly conducted feedback sessions with customers and implemented product surveys in the app.
    - **Regulatory complier:** Set up a weekly alert on common search terms that were relevant to my product and market.
    - **Facilitator:** Met weekly with marketing, sales, customer success, UX research, and QA teams. Met biweekly with leadership. Met monthly with legal and finance.
    - **Data-driven decision maker:** Implemented product analytics and heat mapping to understand user behavior, and updated all internal data products and dashboards after pervasive data QA issues.
    - **Adaptability manager:** Used product strategy sessions with leadership to assess the potential impact of all major pivots.
    - **Conflict resolver:** Reinforced the use of a scoring method when conflicts were too close to call to incorporate an objective influence.

As stated previously in the chapter, you won't always be great at every aspect we touched on for the AI PM role. Building a career as an AI PM is about finding the right combination of company culture, role expectations, and personal preferences for the type of AI PM career. Your personal preferences and priorities will also change as you mature. When I began my career, I was more focused on the technical proficiency aspects of the AI PM role because this was the area I felt I needed the most development on. As I went, I gravitated more toward the leadership and problem-solving aspects of the role because the problems I was solving were becoming increasingly more complex. As I was growing in a leadership capacity, I was also managing more complex teams in larger organizations. As my own aspirations grew, so did my focus and adaptation on my AI PM career.

Anticipate how you'll grow and change in order to build a roadmap for your own career maturity, but remember what makes you you too. Think about the aspects of the role that you really love and feel you're best suited for. While it's nice to be consistently striving to grow and challenge ourselves, I believe our deepest strength comes from knowing ourselves and what we truly value. Only then can we attract opportunities that resonate deeply with who we are and the kind of work we want to do in the world. If I were to think about the aspects of the role that align most with me, my interests, and my skill set, my favorites would be *visionary*, *motivator*, *customer advocate*, *strategist*, and *conflict resolver*.

I feel useful and appreciated when I'm translating a product vision through different manifestations, channels, and discussions and aligning it with the company's mission. It feels meaningful to me to help motivate others back toward a core purpose and focus on what will make the biggest impact on customers and users. I feel a deep sense of purpose when I can successfully resolve conflicts that can keep product teams from doing their best work. As an AI PM, it's important to me that I build a career that makes me feel useful, appreciated, meaningful, and purpose-driven. Find the combination of factors that feels right for you and do your own exercise to understand what values and considerations you want to build into your career trajectory.

Summary

Remember that "good" is in the eye of the beholder. All AI PMs will be subject to the context they find themselves in. Your direct manager, leadership team, corporate culture, and peer relationships will all have an impact on how your work is being seen. Everything we've discussed in this chapter is to be taken as a form of best practice. We can't be all things to all people, and if one were to take care of every aspect discussed in this chapter, they could find themselves in an unsustainable work situation. Rather than taking exhaustive actions to wear every "hat," it's best to think in terms of systems. What systems can you put in place so that you're at least covering some aspect of each AI PM "hat" with some regularity? Remember, it's a hedge, not a guarantee.

Building structure and systems around your work isn't about setting up dogmatic rules, it's about creating healthy boundaries so that we don't do too much and burn out. It's much harder to come back from burnout than it is to avoid it altogether, so be gentle with yourself as much as you can. You're balancing the wills, intentions, frustrations, and conflicts of many people, all at once, in the AI PM role. If you're not careful, you could find yourself with problems coming from all corners. But you can set up guardrails and boundaries in a way that feels supportive. You'll find yourself enjoying all the gifts this career has to offer you for years as a result.

In the next chapter, we will explore what maturity and growth look like in the AI PM career path.

19

Maturing and Growing as an AI PM

If you've been following along through all the chapters of this book and haven't been dissuaded about embracing a career path of an AI PM – congratulations! You've made the decision to embark on an epic journey that will enrich your life. I mean that wholeheartedly. Although I'm not the kind of person who derives their sense of self-worth from the work they do, I find the role of AI PM very rewarding and worthwhile.

My decision to dive into AI and machine learning was a practical one. It's true that I was excited by the possibilities AI could open for all of us. But in pragmatic terms, I chose this field specifically because I thought it would occupy my mind for the next couple of decades. When I started my journey, I didn't know that I would end up becoming a PM. I just knew that I was mystified by the world of AI and, deep down, that nothing is so complicated that you can't understand it if you apply yourself.

I wasn't wrong. Many opportunities opened up for me once I got started, and it's one of the reasons I wrote this book. I wanted to help demystify it for others as best as I could. I found myself getting bored with my previous roles and often thought, "Is this it?" I've never once had that thought as an AI PM and I'm grateful for that. If I ever do get bored, there are many industries and types of products I could expand into. There will always be more to learn and grow from in this career path, and I feel grateful to have chosen something that holds my attention and interest the way this does.

One of the greatest disillusionments I experienced during my first internship in college was realizing that none of the people I was working with in the office liked their jobs. It was just a paycheck, something to do for eight hours a day away from kids and spouses. I didn't want that for my professional life. I wanted to work in a field that kept me constantly learning and growing because, on an atomic level, that was a way to resist the predictability and drudgery of "work." The AI PM career path is for those who like to be constantly challenged and continuously learn. It's best suited for curious, growth-oriented people who value community and collaboration. In this chapter, we will be exploring what it means to mature and grow in the AI PM role. We will cover the following areas:

- Building your ideal AI PM roadmap

- Keeping up with the world of AI and keeping the student mindset alive
- Widening your community and giving back to others
- Deepening your footprint on your professional network
- Anticipating future trends in the AI PM career path

We will be ending with a case study example of practical steps that can be taken by a seasoned AI PM to maintain their place in their field.

Projecting – what's your ideal AI PM roadmap?

As you continue through the process of embracing challenges, reflecting on them, receiving feedback from others, and ultimately sharing your lessons learned with others, you'll start to develop a kind of callus around hardships. They won't affect you as much after you've experienced enough setbacks and emerged triumphant afterward as a result. You won't be as reactive to the experiences around you. You'll be able to then fully and proactively embrace the unknown. The callus that will form from fully embracing a growth mindset will make you professionally resilient and persistent about working toward the kind of career you truly want. This is where the fun starts! Depending on how ambitious you are, you'll start to get a sense of how quickly and in what direction you want to be moving.

Challenges have a lot to teach us, and those who are able to embrace challenges, self-reflect on experiences, seek and incorporate feedback, and share knowledge with others will build enough resilience and understanding of themselves to map out where they're headed in their AI PM career. Building an AI PM roadmap is about exploring evolutions in your professional mindset and your work. Initially, you might be resistant and reactive to challenges, but then you grow into a more curious version of yourself. In more advanced stages, you become more proactive about steering your career and shaping your future more concretely.

In the following sections, we have used levels to establish career milestones, which you can use as suggestions to start to build your own AI PM career roadmap.

Level 1 – building a foundation

In this level, you're building a foundation for your educational background and skills. We covered a lot of these areas in *Chapter 17*. Feel free to refer to that chapter for more context on how to get started on the AI PM career path. This level is all about choosing which learning or education path you want to embark on and how you'll gain your early practical experiences. It shapes your understanding of the field and helps you acquire the essential competencies to advance in it. At this level, that will most likely look like an internship or your first entry-level role. You'll begin the early stages of forming your professional network here as well.

The following are some suggestions to get you started:

- Choose whether you want to go down the formal or online course route to gain technical skills. Think about your budget, time constraints, and personal learning style to understand if you thrive in more structured or more accessible environments.

- Look for internship opportunities with companies that resonate with your values and career aspirations to gain hands-on experience and understand the dynamics of working in a tech environment.
- Consider entry-level positions like product analyst or junior product manager roles, which can offer you a warm introduction to product management or data analysis.
- Join professional groups focused on data, AI, or product management to begin forming your professional community.
- Seek out mentors who have experience in AI product management or related fields who can share experiences and give you guidance on how to move forward in your career.
- Build your professional brand and connect with people on LinkedIn. Ask them questions about their work, daily responsibilities, career path, and perspectives on the industry and where it's headed.
- Embrace this stage with curiosity and determination! The beginning is always hard but your hard work will pay off if you approach it with the right energy.

Level 2 – strategic growth

By Level 2, you're already several years into your AI PM career path. At this point, you have some experience under your belt and are ready to take on more advanced courses or more complex AI products. You've also deepened your PM skills through the roles you've taken on. Because you have more applied PM experience, you're more comfortable with creating strategic plans, roadmaps, market analysis, and user research. At this point, you also have a robust portfolio of case studies, contributions, projects, and products you've supported. You feel more confident in your abilities and you're excited by more complex challenges.

The following are suggestions to ensure that you keep growing:

- Identify areas where your technical knowledge can expand or get deeper, whether that relates to AI/ML models and algorithms, data analytics tools, or AI frameworks that are relevant to the work you're doing.
- If the work you're doing is increasingly specialized in areas like NLP, computer vision, or cloud computing, consider more focused certifications or courses to expand your knowledge.
- Decide if there are any projects you can take on to help you get exposure to deeper technical expertise or give you more experience with data science or ML engineering teams.
- Take inventory of which product management skills you need to strengthen. Refer to the job hats covered in *Chapter 18* to help you assess which aspects of the AI PM role you'd like to improve on.
- As you accumulate experiences, it's important to document all your products/projects/case studies in a portfolio, on LinkedIn, or on other public websites to serve as a testament to your skills, growth, and contributions as an AI PM.
- Keep active with the professional communities you've built and seek opportunities to participate in AI PM forums, workshops, and meetups to connect more deeply with your peers and mentors.

Level 3 – specializing and leading

In Level 3, you've likely specialized in one area of AI or a particular industry. At this level, you've gained familiarity with what it means to manage teams, and the scope of your work is larger. Your focus has shifted toward a culture built around AI and product management within your organization. You may have had exposure to leadership training or developed your soft skills to prepare yourself for potential conflicts and high-level negotiations. You're also playing a more active role in strategic initiatives and publicly engaging with your professional community through talks at conferences or knowledge-sharing opportunities. The impact and influence of your work expand significantly here.

The following are some suggestions for expanding your impact:

- Reflect on whether you're happy in the industry/AI tech focus you're settled into and if your current path aligns with your long-term career aspirations. Of course, you can always make changes, so consider the aspects of your work that bring you the most joy or are most meaningful.
- Take inventory of potential other industries or AI specializations you're excited by. Research trends or go to industry- or specialization-specific conferences and events and connect with others.
- Pursue leadership training opportunities to make you a more effective communicator, negotiator, and conflict resolver.
- Look for leadership training programs or workshops on emotional intelligence, team dynamics, and strategic decision-making to build your tools and skills to lead teams effectively.
- Seek constructive feedback from sources you respect and trust about your effectiveness around larger projects, scopes, and products.
- Identify gaps in your capacity or skill set based on that feedback and use it to inform your professional development plan.
- Lead strategic initiatives in your work environment to make your work more visible and demonstrate your ability to manage complex projects and drive organizational change.
- Find ways to share your work with your professional network and speak at conferences, participate in panel discussions, or write articles for publications you respect where you can.

Level 4 – a light for others

At this level, you've become a master. Your expertise is recognized. You're likely sharing your thoughts and best practices with others often and have created a brand around your AI PM work. You may have even gotten advanced degrees or been published for your work. Perhaps you've even started your own publication. Your focus has shifted to expanding your thought leadership and influence on your AI specialization or industry. You're also likely to speak regularly at large conferences and events and serve as a board member for smaller tech companies. You're heavily involved in your professional community and offer mentorship to the next generation of AI PMs.

Here are a few suggestions for maintaining mastery:

- Identify where you can correct assumptions or misunderstandings people have about the nature of AI products, ethical AI considerations, or the role of an AI PM in published papers, blogs, books, or speaking engagements.
- Look for effective talking points for inspiring junior AI PMs and offering them guidance that truly reflects what's possible for them.
- Share about your journey, the challenges you faced, and strategies that led to your success with your professional community and mentees.
- Consider creating your own mentorship program, workshop, or webinar that's focused on offering practical advice to others navigating their own paths.
- Use your seasoned skill set, knowledge, and expertise to contribute to startups that bring real (and good) value into the world!
- Embrace this stage with a spirit of generosity and a desire to uplift others. You demonstrated devotion and drive to get here. Show others how to bask in their accomplishments. Be a beacon of light for others!

Learning – staying informed and inspired

We covered the importance of upskilling and establishing your technical knowledge when you're just starting out in *Chapter 17*. But once you've matured into your career, you're not "upskilling" anymore, per se. You're maintaining your knowledge and staying relevant in your field. In this section, we will be exploring how to keep learning in a sustainable and inspired way as you get to a more mature place in your AI PM career. The operative word here is *sustainable*. It just so happens that the field is always rapidly changing, which means there's always something you could be doing to stay ahead. Here, we will explore the various ways you can stay up to speed on AI and PM advancements in a way that feels manageable and sustainable so that you can keep being inspired with your work.

Some key areas you'll want to focus on as you mature in your AI PM career are technical skills, product management skills, and soft skills. All these can be expanded on through three main avenues: building your professional brand through thought leadership, bolstering knowledge through certifications and degrees, and sharpening your skills through professional development. In the following sections, we'll examine the various ways you'll be able to learn and stay informed with regard to all three areas to maintain a well-rounded background as an AI PM.

Thought leadership

We all have different interests, values, and skill sets that we bring to the AI PM role. Building a small collection of thought leaders you resonate with strongly is a way to connect with others in your field based on their ideas. It's one of the human experience's most beautiful concepts: the idea that we can feel connected to others by the quality of their ideas alone across time and space. As you mature as an AI PM, you'll be more discerning about the kind of information you let yourself ingest. When you're starting out on your journey, it's harder to tell who's worth following.

There is no shortage of AI, data, machine learning, product management, or leadership influencers and thought leaders online, at universities, or at conferences. Many of them will regurgitate many of the same ideas, hoping to give themselves relevance and notoriety. But as you regularly keep up with AI, product or tech industry newsletters, blogs, articles, and white papers, you'll start to see a few stand out from the crowd. Keep track of the people who inspire you. Make sure they are reputable professionals who you feel truly have something to say.

One AI PM thought leader who inspires me is Mira Murati. Murati has been a vocal advocate for the ethical implications of AI and has emphasized the importance of developers creating systems that are safe, transparent, and aligned with human values. She oversaw the development of powerful AI models like ChatGPT and DALL-E and her public discussions around responsible AI use that highlight AI's risks and societal impacts are admirable, especially when you consider how successful and accessible the products she's brought to market have been. She has held various prestigious positions at Tesla and OpenAI and has shown strength in navigating those dynamics within a predominantly male industry, eventually stepping in as interim CEO of OpenAI during Sam Altman's brief absence. Murati inspires me as an AI leader and practitioner and she's a shining example of what is possible in tech.

The thought leaders you let into your internal world will have an impact on your own philosophies on life and work. They will impact how you make decisions in your work experiences and, ultimately, will have a role in shaping your career. Let your curiosity guide you and make it a habit to stay up to date on their articles and recent work on a weekly or monthly basis. You never know: as you scale up in your career, you could eventually connect, meet, or even work with them. Cover your bases by having preferred thought leaders you connect with across various aspects of the AI PM role: leadership, communication, product management, data, and AI.

Certifications and degrees

As you keep growing and exploring the areas of AI/ML that you want to work in, getting more certifications and coursework could be a great way to fill in gaps you might have in your technical background or your understanding of soft skills. You might even pursue more bootcamps or higher learning degrees if you want to deepen your knowledge or expertise in new areas of AI as well.

I've known an AI PM who left their corporate role, went to pursue a PhD in deep learning, stayed in academia for a few years, and then returned to the private sector to specialize in deep tech startups. While this might sound a bit roundabout, this person had a great time taking a break from the grind to return refreshed. Degrees and specialized certificates can also play a part in pivoting your role as an AI PM. Maybe you want to become an AI PM who focuses more on user-centric design or market strategy. Or perhaps you're sure you want to get involved with the high-level management of AI PM teams. Use these tools in cases where you're looking to make a significant area or focus more prominent in your career trajectory.

> When I think of deepening knowledge, I think of courses and degrees. But when I think about sharpening skills and acquiring tools, I think about professional development. You can think of them as different focus areas with different goals. Professional development is all about workshops, seminars, and conferences – ways you can deepen your professional capacity to be effective in your work. While certifications and degrees focus on deepening your applied, technical skills and credibility through an academic space, professional development is more industry-oriented.

Professional development

You certainly don't have to just keep your skills sharp through bootcamps and degrees. There are more digestible ways available to you of filling skill or domain gaps in your professional background. Maybe you attend a workshop or a seminar to learn from an expert. Maybe you just go to a conference and learn about new areas in the AI PM landscape that you hadn't even considered before. You could even just join a book club. There are also online learning sessions, lectures, and webinars you can enroll in to gain new skills and insights.

As you advance in your career, you could also start to organize these yourself. One of the most advantageous things I did early in my own career was to start to organize events around women in AI. In some cases, it was talks from AI entrepreneurs. In other cases, the events were talks given by female practitioners in the industry. Occasionally, these talks were full working sessions where the speakers would give an overview of a project they worked on, along with a step-by-step tutorial on how they gathered product requirements, data, and feedback to deploy and launch a specific use case of AI.

Investing a work day or the occasional weekend in something like this will have greater downstream effects because your horizons will be expanded each time. They're also a great way of vetting which areas you want to strengthen. Are there certain technical skills you've been curious about but haven't invested much in? Are you curious about leadership but lack the confidence or communication skills to act toward it? Keeping professional development experiences occasional will mean that you'll have plenty of time to integrate the information and insights you gain from these smaller engagements into your own work as an AI PM. Small seeds can have a lot of potential, so use these to start small and see if they inspire you to take more concrete actions for your career.

The following are some practical actions you can take to maintain a learning path that makes sense for you:

- Establish short- and long-term goals. Do you want to learn a new skill set, programming language, tool, or framework? Are you looking for expertise in a particular area of AI? Get clear on what you want to learn and why it will be helpful in your career.
- Establish a sustainable schedule. Are you able to sustain learning on a weekly or monthly basis?
- Clarify which learning paths make sense for you. Do you prefer MOOCs or an advanced degree?
- Find inspiring reading materials. Are there any books, blogs, academic journals, or research teams that speak to you?

Networking – deepening your involvement with the professional community

Engaging with your professional community serves two purposes:

- It allows you to keep learning and growing through the work and insights of others.
- It also allows you to share your own insights, successes, and failures with the very people who can appreciate them the most.

Networking has acquired a bad reputation over the years. It conjures images of stuffy rooms full of strangers, places where you put on a brave face and an air of confidence to "work the room." Maybe you give yourself a quantifiable goal: "Meet five new AI PMs today." Maybe you do indeed speak to them and exchange a few words but nothing meaningful or especially substantial comes of it. You've done the work. You've left your house, met five strangers, and checked off the "networking" box on your to-do list. The connections feel lukewarm at best and the experience is ultimately unfulfilling.

But that's not real networking. That's just personal brand evangelism masquerading as networking. True networking is about establishing connections with people you truly resonate with to form a safety net. That safety net should be based on something real, something authentic. These should be people you can turn to when you feel confused, disheartened, afraid, curious, or inspired. You don't have to become best friends with these people, but you should be able to share some vulnerabilities with them. What's the point of having a professional network if it can't support you when you need guidance or reassurance?

Build strategies that feel natural for you as you start forming authentic relationships with your professional peers and go beyond transactional networking. Start by giving first. Maybe you can offer an endorsement or a recommendation for someone you know without asking for something in return. Or you could host a session or organize an event about sharing successes and lessons learned with colleagues in your organization or network. In a world where people are often approached by those wanting something, demonstrate emotional intelligence by giving something of value before asking. It will set you apart from the crowd and create a foundation for lasting, meaningful professional relationships.

As you gain confidence and start to share more about your work through your own channels, your professional network will keep you honest. Many of them will become your best champions as you start to build a public-facing brand around your work. Here are a few suggestions for how you can get involved in your professional community over time:

- **Online communities:** LinkedIn, Slack, and Discord are powerful platforms that support all kinds of niche groups. An easy way to maintain a professional community is by leaning on these by frequently posting and commenting. Beyond posting and commenting, you can promote your written work, portfolio, or any events you're a part of. You can think of these as the largest part of the funnel. Sure, you may make some deeper connections here over time but many of the connections here will be weaker if you don't find ways to contextualize them in your life.

- **Conferences:** Industry events and known conferences are a great way to synchronously meet a lot of people under one roof. Your chances of meeting people and connecting organically are higher in conferences because you're there in person. The more you go, the more likely you'll see people you've met before. Particularly if you're maintaining communication with people online, you can make sure you're making time to see people you've come to know virtually and root them into the present moment. Conferences are a great way to hear established peers in your industry speak. They also typically have happy hours, afterparties, and musical events after the official schedule, so the chances that you'll connect with people on a more personal level are higher.
- **Meetups:** Regular in-person meetups, particularly in bigger cities, are quite common. With meetups, you get some of the benefits of a conference without the overwhelming tsunami of people. You can receive a lot of the social benefits of seeing people in person without needing to spend nearly as much money. They also offer opportunities for deeper connection because often you're seeing the same faces over and over again. The opportunity to experience belonging is highest with local meetups because you know you're connecting with people who live in your vicinity.
- **Mentorship:** As you expand on networking efforts throughout your career, through the various avenues mentioned above, eventually you'll start to form closer bonds with people. Many of these close bonds will take on a mentorship role. Mentors don't need to be significantly further along in their careers than you. Sometimes they can even be at your level. But they will be people who you can turn to when you're in need of guidance. You won't have the capacity to turn all your professional connections into true mentors, but in essence, the people you let into your inner circle have the capacity to offer you mentorship. Keep the ones you trust and respect close to you and give back to them when you can.

Growing – the student becomes the teacher

It's one thing to learn and grow as a student, but it's another to take that knowledge and guidance and turn it into real wisdom. In many ways, you'll always be learning as an AI PM. But at some point, you will be able to give back. Growth is about taking lived experience and using it to improve as an AI PM and a person. In this section, we'll examine the growth mindset and the areas that contribute most to the application of it.

Embracing challenges

As you continue to mature in your AI PM career path, you'll eventually hit major challenges. Maybe the product you've been working on developing just isn't hitting the performance targets that are needed to meet your customers' and users' expectations. Or even worse, maybe it does but you're still struggling with reaching product market fit or with establishing a consistent user base. Perhaps one of the main takeaways from your market or competitive research deals with the integration of a new form of machine learning you haven't worked with yet for an upcoming feature. Whatever the challenge is, keeping a curious spirit and embracing the challenge is the way through.

I realize that's easier said than done, particularly when you have leadership expectations to balance in the fray of it all to keep your job. The AI PM role is rife with challenges and pressures coming from all sides. Balancing ethical considerations, technical constraints, regulatory compliance, high user expectations, complexity in AI systems, and technical innovations piles a lot on the AI PM role. Complexity in AI systems often introduces unpredictable results, data biases, and issues with interpretability, which adds more difficulty to development and deployment.

But viewing these challenges as opportunities can be transformative. Use them as a catalyst to develop a new skill or improve your problem-solving abilities. The more you do, the more resilient you'll be when they do arise. Embracing a growth mindset means you believe in your ability to excel and succeed through challenges. The idea that persistence, learning, and effort are the way to become more capable is a powerful one. It means that you don't let failures define you. I would venture so far as to say that keeping optimism alive when challenges occur is one of the foundational skills of a PM.

Think about it; you often have to be the one who motivates and inspires the people you work with when they get down. But who can do this for you at your organization? Corporate environments in many cases can be brutal; we often have to become our own cheerleaders. But that doesn't have to be a bad thing. If we never experienced challenges, we'd rest on our laurels and our work would be consistently stellar but dull. The stakes have to be high for a pursuit to be meaningful. Champions are made through struggle. Your ability to dominate that struggle, to not let it consume you or reinforce your inability, is what will keep you growing.

For a powerful example of AI PMs who have mastered the art of embracing challenges, we can look no further than Fei-Fei Li, co-director of the Stanford Human-Centered AI Institute and former chief scientist of AI/ML at Google Cloud. Fei-Fei is an inspiring AI leader who has embraced and overcome significant challenges in AI development. One of the effective ways she found to make an impact was in the healthcare space. By collaborating with hospitals and senior homes in the Bay Area, Fei-Fei found a way to solve the issue of monitoring human behavior with regard to handwashing and hand hygiene, one of the leading causes of infection in hospitals.

Through her work on developing "ambient intelligence" systems, a series of sensors and computer vision models that are able to track the hand hygiene of medical staff, she was able to build AI systems that respond to human needs in real time while minimizing bias that would come from human observers. As a result of her work in hospital settings, Fei-Fei was able to prevent patient deaths and infections, democratize the use of AI in hospital settings, and find a way to repackage some of the same capabilities used in self-driving cars to save lives.

Reflecting

Reflecting is a powerful process that starts with embracing challenges. We don't let challenges define us or send us into downward spirals; we use them as an opportunity to reflect. We can't be afraid to look directly at the shortcomings we experience. This means avoiding the problem, and it will prevent you from ever making the kinds of changes and pivots that will lead to a breakthrough, both personally and professionally. Let yourself be courageous enough to analyze and reflect on the challenges you face so that you can receive their wisdom. Understanding what went wrong and how to prevent it from happening in the future is how growth happens.

On a fundamental level, you're establishing new patterns in your work, honing your product intuition, and building your emotional resilience through your ability to reflect. As you grow and scale in your career, you'll be able to pass down this value and wisdom to the junior members of your team or on to your direct reports if you start to manage PM teams. Reflection turns your hardships and frustrations into actionable next steps. You're not just swimming in a soup of your feelings; you're turning them into a plan. But often, that plan will involve other people, as it should. No one, not even an AI PM, should have to work in isolation. You exist in the context of others. When you're ready to reveal your reflections to others, you're ready for the feedback loop.

Make reflection more actionable as you're navigating your AI PM career, where complex challenges require both technical and strategic introspection. Here are some reflective questions you can explore to help you turn your experience into learning opportunities:

- Which technical skills have I mastered and where do I still feel uncomfortable or need more development?
- How can I prioritize ongoing learning in my career plan?
- Which previous role fit my profile best, based on my technical abilities and product interests as an AI PM?
- How do I want my next role to challenge me in terms of AI/ML or product areas?
- Which AI solutions or products expanded my technical understanding and how can I build on those lessons?
- Do I feel drawn to explore new AI/ML or product areas I haven't experienced yet?
- Which industry problems resonate most with me?
- Are there new verticals or types of AI applications that align more closely with my long-term interests?
- How does my work impact the business, users, or the broader industry?
- Where do I bring the most value as an AI PM?
- Have I cultivated relationships with mentors or industry leaders who can offer me guidance or feedback?
- Which peers or mentors inspire me most and how can I learn from their approach to AI product management?
- Am I comfortable with the leadership aspect of the AI PM role?
- How can I improve guiding cross-functional teams or handling difficult conversations?
- How do I want to use my influence to drive a positive culture around AI for my team or organization?
- What challenges in my AI PM work have tested my resilience most?
- How did I manage these challenges? What did I learn?
- What steps can I take to keep a growth mindset, especially in the face of setbacks and uncertainties in my career path?
- How has my vision for my AI PM career evolved as I've grown?
- What impact do I want to create in the AI field and how can my current career choices bring me closer to that vision?

- What aspect of my work makes me feel most fulfilled?
- What steps can I take to do work that brings me meaning as I advance?

Establishing a feedback loop

In earlier chapters, we discussed the idea of continuous improvement regarding the AI product lifecycle, and this is a principle you can apply to your own growth as an AI PM. On some level, you could argue we are always improving as we reflect on each new job experience. Flexing your growth mindset muscles will keep your strategies adaptive when embracing challenges and help you derive actionable insights from them. This is particularly important in the AI PM career path where products require continuous refinement on the part of the AI PM.

This is because AI products are iterative by nature and require constant improvement to meet evolving user expectations and needs, keep outcomes fair, and maintain relevance in a rapidly changing, competitive, and innovative climate. You will benefit tremendously from keeping an active feedback loop that refines your understanding of these areas. You'll want to seek feedback from a well-rounded variety of mentors, peers, stakeholders, and direct managers to help root those insights in reality. It helps you prevent blind spots in how you approach product challenges, model performance, feature selection, ethical dilemmas, user experience, or balancing technical depth with usability in your own career.

Rooting our growth through the context of others is a way to establish trust with them, and it expands our understanding of ourselves. It's also incredibly inspiring. A person who exudes a sense of impenetrability isn't relatable. False confidence may work for some people for some time, but eventually, reality will come crashing in. Let others see that you're committed to growth and that you won't let your own ego prevent you from improving. Be willing to accept constructive criticism from people you trust and use that as a catalyst for your own personal and professional growth. This is especially relevant in high-stakes, dynamic environments like AI product teams. An AI PM who's open to learning and improving their own gaps reinforces a culture of growth for the whole team and sets a powerful example. It also means an AI PM who's getting better at building safer, more effective, and equitable AI systems and products.

This feedback loop doesn't just have to be one-on-one; it can also be with your greater professional network. As you grow and improve, you'll reach an upper echelon of the AI PM you once were. You will have achieved something that's quite hard for most people to do: you will have evolved. The shedding of your old exoskeleton and the emergence of a new form is something to be shared and celebrated through collaboration. Maybe that means you'll share the challenges you experienced in the form of sharing knowledge about a particular AI use case or AI PM best practices. Maybe the process will open new doors to new mentorship opportunities or valuable connections that you otherwise wouldn't have sought.

Whatever the result is, let it bring you closer to your professional community in whatever form makes the most sense for you. It could come in the form of a blog post, a LinkedIn post, a case study or article, a speaking engagement, or a presentation at a tech conference. Sharing your struggles and lessons learned from navigating the intricate landscape of AI is a real and tangible way you can contribute to your professional community. Even if you think your struggles are common, or that other AI PMs have had the same challenges, remember that no one else is you. Your unique perspective adds value. You never know when the hurdles you encounter could help others who are less far along the mountain. Sharing about your real lived experiences can create powerful connections between people that can potentially lead to new collaborations, work projects, or partnerships in the future.

The following are some strategies and recommendations you can keep in mind as you're building your AI PM career feedback loops:

- Seek feedback from a broad range of stakeholders, mentors, and peers. Make sure to include data scientists, ML engineers, customers/end-users, team leads, and members of your leadership team who you respect and whose opinions you value.
- Keep a record of the feedback you receive and the actions you've taken in response to that feedback. Document the insights and patterns you observe and reflect on how you've applied them in your work or career over time. Let this inform the career narrative you share with others and, most of all, with yourself.
- Prioritize feedback that's specific and actionable. Vague or subjective feedback can leave you with more questions than answers and can cloud your self-reflection as an AI PM.
- Find the balance between confidence in your abilities and humility in recognizing areas of growth. Striving to maintain that balance will enhance your credibility and inspire others to openly engage with you when they struggle with it.
- Treat feedback as a necessary part of your own AI PM career iteration and be genuinely grateful when others share their honesty with you.
- Become an active member of AI and PM communities both online and offline to expand your feedback loop sphere of influence and look for ways to meaningfully add to the discourse.
- Partner with peers or colleagues to co-author articles, white papers, or case studies that address challenges and innovations in the AI PM field.

What's next? The world is our oyster

The future holds a lot of promise for the AI PM career path. As we've gone through this chapter, we've discussed a lot of topics related to growing and maturing as an AI PM. While this is helpful to know, eventually it's less about learning and more about choosing. Each of us is equipped with our own set of values, interests, and world views. As we get deeper into the AI integration of products, we're going to start to see a lot of overlap in AI methodologies. Products will enter the market that employ a variety of AI models in a multidisciplinary way. We'll experience more of a blur between analog reality and AI.

My advice is to try not to deviate from products that you believe *should* exist in the world. Building ethical AI products and systems that are free from bias and harm, bringing forth products that are compliant with regulations and beneficial to the world, is a collective responsibility. Creating AI products and experiences that are human-centric means placing emphasis on user needs, and building AI in a way that's intuitive, reliable, trustworthy, and beneficial has to be the way forward if we're going to build a hopeful future.

As we continue with the AI transformation, we're in desperate need of leaders who will prove themselves to be responsible stewards of AI. As AI PMs, we want to be proud of the work we put forth into the world, particularly at this juncture, when so many legislative efforts are struggling to obtain the footing they need to properly rein in the unintended ill consequences of AI. At its heart, the AI PM role is about strategy, innovation, leadership, and inspiration. This generation of AI PMs is at the forefront of integrating AI into products that redefine how we live and work.

We should use the influence we have to bring AI products to market that enhance human experiences, not diminish them. We have the privilege of addressing some of the most pressing challenges of our time through our work. We can harness the immense opportunity to build trust and bring real value to the world. AI is not just a technology that's relegated to the ivory towers of the tech industry. The AI transformation is everywhere now, across industries, and AI PMs have the unique opportunity to lead this transformation. The world truly is our oyster.

As I expand further into my own AI PM career, I want my next chapter to be focused on helping to build interpretive AI products that create personalized, equitable digital experiences. I want my work to focus on meaningful, empathetic product design that enhances human creativity and productivity. Imagine the possibilities: collaborating in a cross-functional innovation lab to experiment with AI capabilities, participating in AI ethics committees or AI policy shaping. By building AI-native frameworks that balance high-impact features with user-centric values, we can contribute to digital experiences that genuinely improve people's lives. I'm committed to working on products that align with my values and passion, products I can put my whole heart into – what would you like to commit to?

Case study

Astrid is an AI PM with over 20 years of experience in tech and she has successfully led several AI products in the healthcare, finance, and retail industries. She came into AI PM through prior roles in project management and **user experience** (**UX**) and her academic background is business management. Her strengths include a strong understanding of market dynamics, user needs, customer journey, budgeting, and business strategy, as well as leading cross-functional teams. After 20 years in the field, she's more selective about the roles she takes on and she's still learning and growing as an AI PM. Currently, she is the SVP of product at a tech company that specializes in various work optimization tools. In this example, we're going to explore the actions Astrid might take, along with the impact of those actions, as she continues on her path down the AI PM rabbit hole.

Projecting

In the projecting phase of her career, Astrid is looking for ways to make a more meaningful impact on the world around her. She became clear about what kind of work she wanted to be doing as she matured as an AI PM and she was also able to find pathways to bring mentorship and wisdom to the next generation of AI PMs, such as the following:

- **Cross-functional team building:**
    - **Action:** As Astrid rose up at her current organization, she built and led various cross-functional teams. She built a collaborative environment where data scientists, engineers, designers, and domain experts could work on pod teams with a concentration on product suites.
    - **Impact:** This established Astrid as a respected and competent product leader within her organization because it broke down previous silos that existed between development, AI teams, and domain experts. By focusing roles around product suites, these cross-functional teams could build group identities around their areas of focus.
- **Developing a mentorship program:**
    - **Action:** Because her product organization had gotten so big by the time she was SVP, Astrid created a mentorship program to help junior PMs level up their skill sets and expose them to new areas of product they hadn't worked in before.
    - **Impact:** This strengthened her product organization's overall capabilities and collective skill set. It also gave the people who were starting to break into product on her team a way to cultivate and grow in their own career path.

Learning

In the *Learning – staying informed and inspired* section earlier in this chapter, we focused on three main learning areas AI PMs can explore to deepen their knowledge in their field. Thought leadership, certifications and degrees, and professional development are all avenues you can take to keep your learner mindset active as you advance through your AI PM career. Let's explore the actions taken and their impact as Astrid navigates her learning paths:

- **AI upskilling:**
    - **Action:** Astrid had to supplement her AI knowledge by enrolling in advanced AI and machine learning courses through MOOCs like Coursera and edX. Courses included certifications in deep learning, reinforcement learning, and AI ethics. She completes a new one every two years.
    - **Impact:** As AI consistently evolves, Astrid keeps her technical knowledge relevant and fresh through her biannual learning. Even after 10 years, she still feels that there is more for her to learn and keeps it a priority in her professional life.

- **Establishing an AI innovation lab Action:** Given Astrid's 10-year career, she's been able to rise up within her organization and has been leading teams of PMs for the last several years. As a PM manager, she encourages her team to experiment with AI tools they're unfamiliar with and learn new frameworks. She's also created a process of developing MVPs for her team to test ideas from their AI innovation lab quickly.
- **Impact:** This has led to a culture of innovation and experimentation within Astrid's organization and team. It's kept them agile and alert to new AI developments and trends to test out in their lab. It's also created a lot of psychological safety for her PM organization. They trust her and feel she is a co-creator in their professional development and success. This has also encouraged a number of innovative new ideas and solutions that have positively contributed to the company's revenue and user retention targets.

Networking

Getting more comfortable with networking is an effective way of making sure you're growing your ecosystem of peers, mentors, and potential collaborators. Careers need to be maintained with a consistent supply of social opportunities because our passions come alive in collaboration with others. Finding a way to network in a way that feels meaningful and authentic is a fine balance, but when done right, it's well worth the effort. Let's explore how Astrid found opportunities to network and expand her professional sphere:

- **Remaining active in the AI conference circuit**:
    - **Action:** Astrid regularly attends AI conferences and has created a community of industry professionals she sees at conferences.
    - **Impact:** This has given her a strong sense of community and belonging within her career path. It's helped her keep imposter syndrome at bay and it's routinely inspired her as she catches up with people she's known for over eight years now. In hard labor markets, this has contributed to a safety net for Astrid. She's found multiple work opportunities over the years through her chosen AI PM family because she's built trust with the people she's met, consistently been present with others, built organic connections over time, and even spoken at these conferences several times.
- **Building inspired partnerships through academic collaboration:**
    - **Action:** One day, Astrid went to an AI talk at one of her local universities and connected with a woman who served as head of professional partnerships at the university. Initially, they spoke about an internship program, but this led to the development of a strategic partnership where her AI innovation lab could collaborate with research teams at the university for special joint projects.
    - **Impact:** This gave the AI innovation lab at her organization the jet fuel they needed to discover new, unique ways of applying AI in a variety of product use cases. It created a funnel of new, emergent junior talent that was a cultural fit. The partnership also provided additional AI resources for complex projects when the engineering team needed it.

Growing

In the growing phase of your career maturity, you're more focused on expanding your awareness. You look for opportunities to deepen your understanding of your AI specialization, how you want to grow professionally, and where there might be blind spots. Embracing challenges, reflecting, and expanding your feedback loops are all ways you can delve more fully into your own career. Astrid was able to deepen her involvement with the startup world and discover uncharted areas in AI ethics through the following actions:

- **Working on a healthcare startup MVP:**
  - **Action:** Astrid helped ideate, developed, and shipped a predictive analytics tool for a healthcare startup that leveraged NLP to predict patient readmissions. She had never worked in the healthcare domain before this product and took on the role after going to an AI entrepreneurship happy hour where she met a startup founder who was looking for someone with her AI background.
  - **Impact:** After the product was launched, it was implemented in three of the hospitals in her area. She was able to positively impact her local community, improve patient outcomes for real people, and showcase her ability to apply her knowledge and expertise in a new domain space.
- **Setting ethical AI standards:**
  - **Action:** Astrid spearheaded a new ethical AI standard of best practices for the product management, data science, machine learning, and development teams to follow when they're working on AI products within her organization. The rollout also had an extension for non-technical teams to be able to flag AI ethics conflicts even when they're not directly involved in a technical capacity. The guidelines she developed dealt with issues of fairness, bias, transparency, and accountability for all AI products.
  - **Impact:** Astrid was able to mitigate risks that could compromise her company's brand through her investment in the ethical AI standard. She was able to support her team by creating infrastructure around ethics so that no one person felt solely responsible for ethical AI oversight. She was also able to make AI and its potential dangers more accessible to those who were non-technical within her organization. Model performance and fairness improved slightly for the AI products that were already released, which also enhanced customer and user trust.

What's next?

Astrid's next major step is to find a way to combine many of the elements she's touched on in various areas of her career into an entrepreneurial venture that can be her very own. Because of her work with the AI innovation lab, AI conference participation, and academic partnerships, Astrid decides to launch her own AI consultancy business. She realizes she already has a lot of contacts through the various talks and sessions she's given over the years.

She's also worked collaboratively with academics and with the teams that participated in the lab. Astrid's seen her fair share of project diversity and has learned a lot from building ethical AI standards at her own company.

Because of her passion for improving patient outcomes, Astrid wants her consultancy business to focus on the healthcare sector. Working on the NLP-based MVP gave her a sense of confidence in her work, and she realizes there are many startups that she could help like this. Sometimes one experience can ignite a newfound sense of purpose. Working with startups allowed Astrid to open her horizons to new ways of empowering teams. She was used to more rigid, corporate environments prior to the MVP and hadn't considered how valuable her work would be when applied to startups that could really benefit from her many years of experience.

Overall, these various actions and impacts contributed to Astrid's effectiveness as an AI PM leader in her field. First and foremost, it kept her engaged with learning and growing as a practitioner. But these collective actions also created pathways for others below her to learn and grow as well. She was able to keep her teams agile and inspired on a regular basis. Cross-functional teams learned to trust each other and this resulted in more inspired solutions to business problems. It also made everyone, including Astrid, more resilient. Astrid's commitment to continuous learning, innovative leadership, ethical AI practices, team management, and strategy made her an AI PM success story. As a result, she's been able to stay relevant and competitive in her field.

Summary

We began this chapter by covering what growth looks like as you mature through various levels of AI PM maturity and saw a few different areas of learning as you advance through these stages. We covered the topic of effective networking and the role it plays in keeping a long career recharged. We finished with practices and habits that can provide sources of growth in the advanced stages.

As we end this book, remember that reaching maturity in an AI PM career path isn't just about becoming adept at the technicalities of AI products – from ideation to deployment; it's also about developing the resilience, vision, and adaptability that will define you as a true leader in this evolving space. There will always be an ongoing dance between continuously sharpening your technical acumen and cultivating the skills you need to support successful and ethical product launches. The more you progress, the more challenges will feel less like barriers and more like pathways to deeper insights about your work and yourself.

20

Unlock Your Book's Exclusive Benefits

Your copy of *AI Product Manager's Handbook, Second Edition* comes with the following exclusive benefits:

- Next-gen Packt Reader
- AI assistant (beta)
- DRM-free PDF/ePub downloads

Use the following guide to unlock them if you haven't already. The process takes just a few minutes and needs to be done only once.

How to unlock these benefits in three easy steps

Step 1

Have your purchase invoice for this book ready, as you'll need it in *Step 3*. If you received a physical invoice, scan it on your phone and have it ready as either a PDF, JPG, or PNG.

For more help on finding your invoice, visit `https://www.packtpub.com/unlock-benefits/help`.

> **Note**: Bought this book directly from Packt? You don't need an invoice. After completing *Step 2*, you can jump straight to your exclusive content.

Step 2

Scan the following QR code or visit `https://www.packtpub.com/unlock/9781835882849`:

Step 3

Sign in to your Packt account or create a new one for free. Once you're logged in, upload your invoice. It can be in PDF, PNG, or JPG format and must be no larger than 10 MB. Follow the rest of the instructions on the screen to complete the process.

Need help?

If you get stuck and need help, visit `https://www.packtpub.com/unlock-benefits/help` for a detailed FAQ on how to find your invoices and more. The following QR code will take you to the help page directly:

Note: If you are still facing issues, reach out to `customercare@packt.com`.

packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

At www.packt.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:

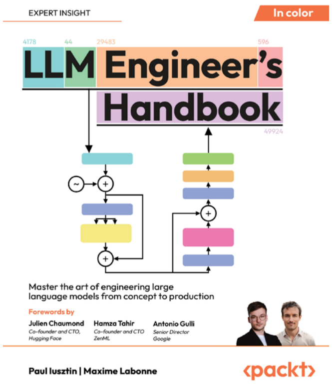

LLM Engineer's Handbook

Paul Iusztin, Maxime Labonne

ISBN: 9781836200079

- Implement robust data pipelines and manage LLM training cycles
- Create your own LLM and refine it with the help of hands-on examples
- Get started with LLMOps by diving into core MLOps principles such as orchestrators and prompt monitoring
- Perform supervised fine-tuning and LLM evaluation
- Deploy end-to-end LLM solutions using AWS and other tools
- Design scalable and modularLLM systems
- Learn about RAG applications by building a feature and inference pipeline

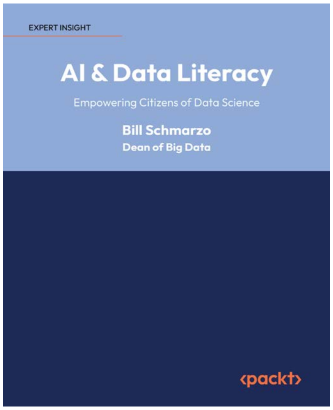

AI Data Literacy

Bill Schmarzo

ISBN: 9781835083505

- Get to know the fundamentals of data literacy, privacy, and analytics
- Find out what makes AI tick and the role of the AI utility function
- Make informed decisions using prominent decision-making frameworks
- Understand relevant statistics and probability concepts
- Create new sources of value by leveraging and applying AI and data
- Apply ethical parameters to AI development with real-world examples
- Find out how to get the most out of ChatGPT and its peers

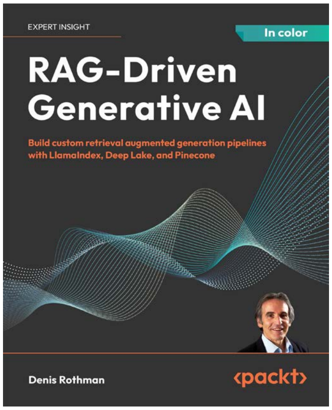

RAG-Driven Generative AI

Denis Rothman

ISBN: 9781836200918

- Scale RAG pipelines to handle large datasets efficiently
- Employ techniques that minimize hallucinations and ensure accurate responses
- Implement indexing techniques to improve AI accuracy with traceable and transparent outputs
- Customize and scale RAG-driven generative AI systems across domains
- Find out how to use Deep Lake and Pinecone for efficient and fast data retrieval
- Control and build robust generative AI systems grounded in real-world data
- Combine text and image data for richer, more informative AI responses

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit `authors.packtpub.com` and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Index

A

B

C

D

L

M

N

S

T

U

V

W

X