

The AI Value Playbook

How to make AI work in the real world



Lisa Weaver-Lambert

⟨packt⟩

THE AI VALUE PLAYBOOK

HOW TO MAKE AI WORK IN THE REAL WORLD

LISA WEAVER-LAMBERT



In Praise of The AI Value Playbook

“Unlock the secrets of rapid AI implementation with this indispensable guide by Lisa Weaver-Lambert. Drawing from expert insights, this essential resource reveals how to harness storytelling and alignment, create an engaging environment for communication, build trust, invest in continual learning, and focus on strategic alignment. With a writing style that’s both concise and compelling, Lisa shows you how to turn AI into a game-changer for your organization – not just a buzzword. Dive in and discover the formula for success.”

— John Winn, Managing Director at Blackstone

“I cannot think of another source where so many AI examples are compiled into one book. I think this is a must read for anyone who wants to learn the history, creation, and, more importantly, ideas of AI to apply to the reader’s own journey. As a leader in the technology industry, I’ve seen firsthand the transformative potential of AI, but also the critical importance of using it responsibly and securely. The AI Value Playbook offers a roadmap to maximize its potential. We can only adapt and change when we listen, and in this book there are many voices across geographies and industries. I highly recommend it.”

— Kelly Bissell, Corporate Vice President, Microsoft

“The AI Value Playbook is valuable for organisational leaders who are seeking to improve productivity by driving value from data and AI. The AI Playbook offers a powerful guide to help companies master AI concepts, overcome the inevitable obstacles, and learn from real-world examples of applied AI.”

— Soumitra Dutta, Peter Moores Dean and Professor of Management, Saïd Business School, University of Oxford

“In an era where data is the lifeblood of success, The AI Value Playbook stands out as the essential guide for companies navigating the transition to an AI operating model. Busy users can leverage chapters for quick reference, while even advanced users will find it valuable as a source of clear explanations for colleagues and friends. Lisa unpacks success strategies, weaving insights with captivating case studies that provide executives with actionable mental models. The book offers a structured approach for initiating, measuring, and adapting AI projects, acknowledging the iterative nature of ROI. Executives seeking AI-driven growth will find it a must-read, with practical guidance and a strategic framework for implementation, including the critical development of talent. It breaks the mold of the typical executive-focused reads. It’s an indispensable ally for business leaders that will leave you inspired and informed.”

— Bogdan Crivat, Corporate Vice President, Azure Data Analytics, Microsoft

“At Race Capital, we invest in transformative technologies. The AI Value Playbook isn’t just about technology. It’s about empowering people. You’ll discover how to integrate AI to future-proof your business and empower your workforce. The AI Value Playbook is your practical guide and roadmap to unlocking the power of AI for your business inspired by real-world case studies.”

— Alfred Chuang, General Partner at Race Capital, Silicon Valley’s most passionate CEO Coach, Co-Founder, CEO, and Chairman of BEA Systems

“The AI Value Playbook offers a clear-eyed view of how AI is transforming industries. Lisa’s actionable framework, combined with real-world case studies, empowers leaders to navigate the complexities of AI integration and release its potential for sustainable growth. A great read for business leaders looking to develop a strategic roadmap to leverage AI for increased efficiency and competitiveness.”

— Pegah Ebrahimi, Co-Founder and Managing Partner at FPV Ventures

“In today’s rapidly evolving landscape, AI is no longer a future possibility, but a present necessity. The AI Value Playbook offers a clear roadmap for business leaders to navigate the exciting world of AI and transform it into a strategic advantage for our organizations. Lisa’s insights and real-world examples across industries and geographies are invaluable for anyone looking for practical techniques to implement AI and drive positive results. If you’ve questioned your company’s ability to innovate - or stay relevant - in an era of exponential technological growth, this book is a must-read.”

— Guglielmo Angelozzi, CEO, Lottomatica Group

“The AI Value Playbook is a roadmap for CEOs to leverage AI to future-proof their organization. Lisa combines insightful analysis with practical case studies, demonstrating the importance of data management and how AI can drive efficiency, accelerate innovation, and unlock new markets. This book is a critical resource for any organization looking to scale with AI and secure a competitive advantage in the years to come.”

— Sajjad Jaffer, GrowthCurve Capital; Ex Co-Founder, Two Six Capital

“The AI Value Playbook equips board members with the knowledge, practical tools, and competitive edge needed to provide the vital leadership required to successfully navigate the transformative potential of AI. By providing insights from real-world case studies spanning hugely diverse organizations, the book reveals how AI can disrupt and redefine industries, driving increased productivity, cost optimization, and long-term growth. It also tackles the gap between AI’s promise and its practical impact, addressing internal organizational barriers. Most importantly, the book equips boards with a practical framework distilled from practitioners that have the bruises and scars from deploying AI in businesses, guiding board members to ask the right questions of their organizations throughout the AI transformation.”

— Sanjeevan Bala, Independent Non-executive Director (Bakkavor, SThree)
| Group Chief Data & AI Officer | AI Board Advisor

“In the fast-paced world of private markets in which I work, staying ahead of the curve is paramount. The AI Value Playbook offers a clear-eyed view of how AI is transforming industries. Lisa’s actionable framework, combined with real-world case studies provides a much needed personal, story-telling approach to a topic that’s understandably exciting yet for many, can feel overwhelming to truly understand. Her deeply humanistic and practical guide empowers business leaders to identify high-potential AI applications. This book is a valuable guide on the challenges and transformative power of AI that speaks directly to the non-technical business leader.”

— Alice Murray, Founder, Editor-in-Chief, The Line. Head of Content and Communications, Alpha Group

“Lisa is an ex-colleague of mine at Microsoft and I respect her opinions. Also, I suspect like a lot of you, I have mixed feelings about the current AI boom: I can see the value in AI but I can also see the vast amount of hype and the obviously ridiculous claims being made. More than anything I see senior executives talking confidently about a subject I’m sure they don’t understand, and that is clearly a big problem. This book aims to help solve that problem by providing a practical guide to AI for non-technical leaders, in the form of a series of case studies and interviews with entrepreneurs and C-level people in the AI space. This is a very readable book – Lisa has talked to a lot of interesting, knowledgeable people – and the format makes it a lot more palatable for the target audience of your boss’s boss’s boss than your average tech book. As a technical person who isn’t by any means an AI expert I also enjoyed reading it.”

— Chris Webb, Principal Program Manager at Microsoft

The AI Value Playbook

Copyright © 2024 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Group Product Manager: Niranjan Naikwadi

Publishing Product Manager: Nitin Nainani

Book Project Manager: Shambhavi Mishra

Senior Editor: David Sugarman

Copy Editor: Safis Editing

Indexer: Manju Arasan

Production Designer: Vijay Kamble

Senior DevRel Marketing Coordinator: Vinishka Kalra

First published: August 2024

Production reference: 4310125

Published by Packt Publishing Ltd.

Grosvenor House

11 St Paul's Square

Birmingham

B3 1RB, UK

ISBN 978-1-83546-175-4

www.packtpub.com

About the Author

Lisa Weaver-Lambert has been leveraging data to solve complex commercial challenges for over two decades through her roles at Microsoft, Accenture, and in capital and public markets. In addition, she has held executive line management positions and served at board level. Her work in regions including Europe, the US, APAC, and Africa has brought her a global perspective.

Lisa's early career in financial services, working on acquisition, retention, and fraud detection programs, provided a valuable foundation in data-driven decision making. This experience solidified her understanding of data quality, governance, and its commercial potential. She witnessed firsthand the power of large datasets and how working with data could solve real-world business problems.

Since then, she has worked with some of the world's best-known brands towards success, earning her recognition as a leading woman in technology in the EMEA investment field. Her extensive experience in strategy and business transformation has led her to collaborate with organizations of different sizes and maturity levels to accelerate performance.

Recognizing the transformative potential of AI, Lisa's mission is to equip non-technical leaders with the knowledge and confidence to leverage AI for their organizations. Lisa recently founded Oxford AI Studio. This initiative aims at bridging industry with researchers and knowledge institutes to tackle real-world business challenges and drive the development and application of responsible AI innovation.

Lisa earned her degrees from the University of Lancaster and ESCP Business School in Paris, strengthening her qualifications through executive programs at MIT Sloan School of Management and UC Berkeley School of Information.

LinkedIn: [lisa-weaver-lambert](#)

Acknowledgements

This book, like the discussions it seeks to facilitate, is the culmination of countless conversations and collaborations. I am deeply grateful to all the interviewees, and their teams, whose insights, stories, and lessons are the essence, of this work. Thank you for your time and generosity of spirit.

I am deeply grateful to the following colleagues and friends who have enriched this book in invaluable ways:

Giuseppe Saltini brought his twenty-five years of experience in the data and information architecture space. Giuseppe's insightful feedback and edits not only enhanced the manuscript's accuracy but, I trust, will also make it more relevant to readers.

David Goodchild lent his creative expertise to the book's visual identity. From the captivating cover design to the clear internal layout and informative diagrams, David's contribution significantly improved the overall presentation.

Kate Willis for her editorial contributions to the interview drafts. And Lara Lambert for her meticulous proof-reading.

Insight Partners for their expertise in identifying high-growth generative AI companies.

The Packt team, whose keen eyes ensure the book's clarity and professionalism.

About the Reviewer

Tanmaya Gaur is a Principal Architect at T-Mobile US, Inc., with over 15 years of experience building enterprise systems. He is a technical expert in the architecture, development, and deployment of advanced software and infrastructure for enhanced user support in telecom applications. He is passionate about utilizing composable architecture strategies to aid creation and management of reusable components across the entire spectrum of web development, from front-end UX to backend coordination and content management. His current focus is building Telecom CRM tools that leverage micro-front-end, artificial intelligence, and machine learning algorithms to boost efficiency and improve the product experience.

Contents

Introduction	1
Overview of AI Concepts and Technology Stack	7
Sam Liang, CEO of Otter.ai	25
Amr Awadallah, Founder and CEO at Vectara	35
Philipp Heltewig, Co-Founder and CEO at Cognigy	49
Miao Song, Chief Information Officer at GLP	65
Ruben Ortega, General Partner at Enjoy the Work	73
Joshua Rubin, Principal AI Scientist at Fiddler AI	83
Nadine Thomson, Global Chief Technology Officer at GroupM (WPP)	99
Sarvarth Misra, Co-Founder and CEO of ContractPodAi	109
Edward Fine, AI and Data Science Consultant, Technologist, and Instructor	123
Sanjeevan Bala, Group Chief Data and AI Officer at ITV	131
Nathalie Gaveau, AI Tech Entrepreneur and Board Member	145
Phil Harvey, Applied AI Architect	155
Elizabeth Ajayi, Director, Intelligent Industry at Capgemini Invent	167
Louis DiCesari, Head of Data, Analytics and AI at Levi Strauss & Co.	173
Vickey Rodrigues, CTO/CDO in Insuretech, Payments, and Healthtech	181
Sean McDonald, Former Global Chief Innovation Officer at McCann Worldgroup	189
Julie Gray, Head of Data and Internal Systems at Agilio	199
Peter Jackson, Chief Data and Technology Officer at Outra	209

CONTENTS

Mark Beckwith, Director of Data Governance and Architecture at the Financial Times	221
Kshira Saagar, Data Science at Wolt and Doordash	229
Joe Romata, Global Head of Customer Experience at a Multinational Energy Company	239
Tomasz Ullman, Former Global Head of Data Science and Strategy at Ford Pro	247
Oz Krakowski, Chief Business Development Officer at Deepdub	257
Case Study: LLMs and RAG Enable Hyper-Personalized Education for Healthcare Technicians	267
Case Study: AI Personalization Increases Engagement in Nascent Tech Communities	285
Case Study: AI-Powered Virtual Agent Augments Service Efficiency	301
Case Study: Generative AI Creates a Paradigm Shift in Innovation Processes	313
Case Study: Unlocking Profit Potential – Leveraging Enterprise Data for Customer Profitability	331
Case Study: Minimizing Customer Churn with AI	343
Case Study: Enhancing Marketing Strategies Through the Power of LLMs ..	365
Case Study: Multimodal LLMs Redefine Software Development and Customer Innovation	379
Where We've Got To and What's Next	395

Author's Note

Writing about AI for non-technical executives and board members requires striking a balance between rigor and readability, making sure to be succinct without sacrificing nuance. I have drawn on my own experience working with leadership teams to find the sweet spot on the continuum, getting the substance right while keeping the book accessible.

Opening up my personal access to leading minds in today's business landscape has given me a unique perspective. Through our compelling conversations, I have aimed to bridge the gap between the technical world of AI and the strategic considerations of leadership. My interview approach maintains the authenticity of the dialogue. I hope you find that I have achieved this and these insights inspire you.

1

Introduction

Sitting across from an astute CEO I had been working closely with, I could sense his growing understanding of AI's potential to accelerate the growth of his business. His plan to integrate two market-leading businesses was complex yet also a significant opportunity. His questions were increasingly insightful, delving deeper into his organization's processes and people. He was envisioning the future – one powered by AI.

It was a chance to rethink the operating model of his company. As the weeks passed, he began networking with other external leaders, who were all harnessing data as a strategic asset, leveraging it to gain insights, make informed decisions, and drive growth.

But, as I witnessed this shift, new questions began to form: How could I scale this understanding through human stories? How could these practices and learnings be applied to other organizations? This marked the beginning of *The AI Value Playbook*, a practical guide to AI, aimed at supporting non-technical executives and board members to quickly formulate a perspective on how to integrate AI into their businesses. AI is no longer new, but extracting value from it remains a challenge for many of us.

INTRODUCTION

Driving business value with AI

Today, the expectation that AI will inevitably create value, foster innovation, and secure competitive advantage is prevalent in everyday life. Our daily interactions with products and tools from leading technology companies – Amazon, Apple, Facebook, Google, Microsoft, and Netflix – create an assumption that AI will flick a switch and deliver results almost instantly.

But this is not the reality for most organizations. The tech giants have been built from the bottom up with data as their bedrock and with AI integrated into core processes for nearly a decade. Technological innovation is intrinsic to their business models and growth strategies, and they face few of the cash, storage, or resource challenges common to others.

However, the transformative potential of AI is not exclusive to these tech giants. AI promises to redefine businesses, scaling across every sector of our economy. There's clear evidence that integrating AI into business operations measurably enhances efficiency while reducing costs. It is now among the top five value drivers on CEO agendas.

The growing cost of capital, unstable inflation, and pressure on operating expenses are all compounding stresses on organizations today. Therefore, the ability to manage and control processes underpins growth, regardless of the sector. Goldman Sachs' research asserts that AI has the potential to boost annual global GDP by 7% (or almost \$7 trillion) over the coming decade (packt.link/f5sli). Therefore, the transformative potential of AI promises to boost productivity and fuel long-term economic expansion.

In light of aging populations and other economic growth headwinds, productivity becomes even more critical. Without a sustained rise in productivity, economies will face a growth ceiling that will impact all businesses.

Sustained productivity growth therefore depends, in part, on how effectively AI solutions are integrated across sectors. However, the speed and scale of this integration remain uncertain. We're at the early stages of this journey.

INTRODUCTION

Where we are

History teaches us that there are inherent delays with any transformative technology. The Solow paradox, coined by economist Robert Solow in 1987, claimed that while you could see computers everywhere, their influence was not then evident in productivity statistics. This paradox was resolved in the 1990s when large sectors of the economy, such as retail and wholesale, finally adopted technologies such as client-server architectures and ERP (enterprise resource planning) systems. These innovations helped revamp business processes, particularly in supply chain and distribution center efficiency, leading to noticeable gains in productivity.

Today seems to mirror the patterns of the past. Despite the fact that AI is ubiquitous, it hasn't yet translated into significant productivity growth. This discrepancy exists primarily because large sectors of our economy have yet to integrate AI into their businesses. While the tech sector is hinging its future strategies on the continued development and application of AI, many other sectors are still navigating their path.

Governments worldwide are turning to AI as a potential engine for both economic growth and competitive advantage. Discussions at this level frequently revolve around innovation, regulation, talent, and specifically the systemic barriers hindering progress.

As technological advancements have brought down costs, the financial hurdles associated with AI adoption have significantly reduced. So, as cost and access are no longer obstacles, the systemic challenge lies in developing AI capabilities within organizations and strengthening the integration muscle. This is a process that demands time and commitment.

The crucial question remains – what's the optimal way for businesses to develop and integrate AI?

INTRODUCTION

Who this book is for

This book is a practical guide to AI, aimed at supporting non-technical executives and board members to quickly formulate a perspective on how to integrate AI into their businesses. It includes which levers and processes to consider to future-proof value.

The emphasis is on questions that are frequently posed, such as, “How can businesses adapt to these emerging technologies? How can they start building and deploying AI as a strategic asset to drive efficiency? What risks or threats need to be considered? How quickly can value be created?”

This book is a response to those demands. In a series of in-depth and wide-ranging conversations with an elite group of practitioners, from CEOs leading new generative AI-based companies to data scientists and CFOs working in more traditional companies, we hear how they have succeeded in building AI solutions in businesses of all sizes across the globe.

Our focus has been on how to integrate AI into organizations to drive value. However, as businesses are at varying levels of maturity, the conversations in this book cover a variety of applications from cutting-edge generative AI to the essential fundamentals of data management.

Real-world conversations

A diverse scope of interviews and case studies provide honest reflections and insights from their protagonists. We discover the reasons behind their choices, the obstacles they’ve faced, how they’ve persevered, and the lessons learned from their own set of experiences. Their thoughts and conclusions deserve attention and make for inspiring reading.

INTRODUCTION

I have distilled core principles and actionable strategies in a practical playbook. Yet, it goes without saying that not every tool or practice is universally effective; what works for one organization might not be effective for another. However, there are consistent principles that can be learned from others, even if the details will differ within each organization. The future application of AI technologies is likely to be an amalgamation of different models and practices, with the deployment of solutions that respond to evolving business circumstances and needs.

Just as a map is a description of an environment, but doesn't tell you where to go, *The AI Value Playbook* provides a roadmap to navigate your route forward based on the stories of individuals who have made meaningful progress.

Starting to work with AI doesn't mean immediately overhauling your current operating model or handing everything over to data scientists. Instead, it implies working in alignment with your business objectives and key challenges and identifying those that could be better solved with AI. This involves making business opportunities visible through method, through process, and the work done in multi-disciplinary skilled teams. Incorporating AI is fundamentally about empowering people to do their best work and make the best next decision.

If you are evaluating how to build AI capabilities to future-proof your company, this book will provide you with the necessary guidance. Through insightful interviews, real-life case studies, and actionable strategies, *The AI Value Playbook* will help you better organize and harness the value of AI.

How this book is organized

As a versatile resource, this book can serve as your AI companion, relevant to wherever you are on your AI journey. It allows you to access sections or examples that directly address the specific challenges you face.

INTRODUCTION

It's structured into four distinct sections for easy navigation:

1. An overview of AI concepts and technology stack provides relevant context to engage with and learn from for the subsequent conversations and case studies.
2. In-depth conversations with experienced practitioners and leaders across various sectors, sharing insights on their data and AI journey. These dialogues are designed to be read in any order and offer a broad perspective on the implementation and challenges of using AI. Some businesses are mature in their use, others are just beginning their journey, and many are at various stages in between.
3. Case studies exploring the specifics of real-world applications. These present detailed analyses of practical scenarios, offering a closer look at the application and impact of AI.
4. *The AI Value Playbook* is the practical framework I have distilled from the people I've spoken with and my own professional experience for successful AI implementation.

2

Overview of AI Concepts and Technology Stack

To give context to the interviews with leaders and practitioners that follow in this book, it is important to set a foundational understanding of some of the fundamental concepts and taxonomy of AI. However, AI terminology is challenging today due to its dynamic nature which is resulting in new terms and redefinitions. Therefore, the focus here is on what terms will be beneficial to the reader looking to understand how businesses are using AI to drive value. This chapter will help you engage with and learn from the interviews and case studies.

This chapter is presented in two key sections:

- **2.a AI concepts:** This section delves into the core principles that underpin AI and the different types of AI, such as machine learning, deep learning, neural networks and generative AI. It also highlights the strengths and limitations of AI technologies, helping you identify which ones are most relevant to your business needs.
- **2.b AI technology stack:** This section provides a brief overview of the set of technologies, tools and services used together to build and run an AI application, establishing and ensuring trustworthy AI Systems and AI as a Service (AIaaS).

2.a AI Concepts

The term “artificial intelligence” has been used in computer science since the 1950s, with the Dartmouth conference in 1956 usually cited as its official birthplace.

Since then, the field of AI has been through a series of cycles of strong support and investment versus periods of criticism and reduction of investment, referred to as AI winters. These cycles throughout the ‘70s, ‘80s, and ‘90s had a material impact on consistent progress.

This situation fundamentally changed with the advent of big data and cloud computing. The affordability and power of data storage and CPU time increased significantly, making it possible to train models more effectively. As a result, AI models that were once theoretical and unattainable due to resource constraints became practical and efficient.

In the last five years, the largest technology organizations in the world have tackled and started to solve problems that were completely intractable only ten years before. We are at the start of a period of high investment and rapid progress that is bringing new enthusiasm among a larger population. However, the outcome of this period remains uncertain.

AI is a broad definition for a range of software applications. As a branch of computer science, AI can be thought of as a smart computer that can do tasks usually done by people. It understands languages, recognizes patterns, solves problems, makes decisions, and learns from experience. It uses vast amounts of data and follows instructions called algorithms to automate tasks.

Comparing AI to traditional software applications, for instance, if we want software to detect dogs in a picture, writing conventional software becomes complex and costly. AI is more suitable for this task. The process involves feeding an AI system with thousands of pictures, some containing dogs and others not. The AI learns from these examples and develops an internal configuration that can detect if a random image contains a dog. This example illustrates how AI is suitable for solving complex tasks that are near impossible to tackle with

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

traditional software. This is a fundamental shift from traditional programming which is designed to produce deterministic results based on a set of inputs, as AI learns from the data and makes predictions or decisions without having to be explicitly programmed to perform the task.

Types of AI

At a high level, AI can be categorized into two main types: narrow AI and general AI.

- Narrow AI, also known as weak AI, is designed to perform a specific task, such as voice recognition or image analysis. These solutions operate under a limited set of constraints within a limited context and are only “intelligent” within their specific domain. Examples are frequently found in finance, for example, scaling commercial contract reviews, augmenting collection prioritization, forecasting, streamlining invoice approvals, and anomaly detection. The examples given in the interviews in this book provide numerous use case demonstrations of Narrow AI.
- General AI, also known as strong AI and Artificial General Intelligence (AGI), refers to systems that possess the ability to perform any intellectual task that a human can do. They can understand, learn, adapt, and implement knowledge in different domains. While this type of AI is not currently available, prominent AI researchers forecast that the first AGI will be developed within the next decade. Nevertheless, it is the core of many philosophical debates as while the advancement of AI may be inevitable, its ultimate destination is not clear. AI’s function is complex and uncertain. Some Large Language Models (LLMs) have demonstrated advanced capabilities beyond narrow AI; however, there are still substantial differences between these AI systems and human cognitive systems.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK


AI encompasses various interconnected areas:

- Machine Learning (ML) is a subset of AI. It is the foundation that includes the algorithms and statistical models that underpin AI. It refers to the approaches associated with getting computers to learn and act like humans, improving their learning over time in an autonomous fashion, by feeding them data and information in the form of observations and real-world interactions.
- Neural Networks (NNs) are a subset of machine learning, modeled after the human brain's architecture and designed to emulate its processing style. These networks consist of interconnected nodes, analogous to the brain's neurons, organized in layers. Data enters through the input layer, is processed through one or several layers where computation and transformation occur, and culminates in an output layer that delivers the final results. Each layer's nodes are linked with weights – numerical parameters that are adjusted during training – enabling the network to discern complex patterns and “learn” from data.
- Deep Learning (DL) is an advanced subset of NNs characterized by its use of several hidden layers that add complexity and the ability to perform complex tasks. Deep learning approaches excel at processing and learning from extensive datasets and are instrumental in fueling many of the latest innovations in AI. These include generative AI models, computer vision, and self-driving software. The advantage here is the ability to handle and learn from very large datasets, allowing it to recognize complex patterns and make intelligent decisions. The drawback is the almost complete absence of explainability. For example, it does not explain *why* a dog is not recognized in a certain picture when there clearly is a dog present, or why the AI has denied credit to a specific person.
- Generative AI (GenAI) is a subset of deep learning that generates multi-modal content such as text, videos, and images with user interactions. It is trained on vast datasets to detect patterns and creates outputs without instruction. Unlike databases, these models don't query and return 100% accurate historical information. Instead, they mimic inputs and provide new content that is likely to be correct based on probabilistic techniques.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

AI training techniques

Training a custom model frequently involves the use of multiple techniques, as shown in the following diagram.



Note: The above picture is not an exhaustive list, but a representation list of ML techniques used most frequently to solve business problems

Figure 2.1 – A variety of AI training techniques

Some common techniques include:

- Supervised Learning, a traditional approach in machine learning where a model is trained on a labeled dataset. In this context, “labeled” means that the dataset contains the answers to the questions we want the model to be able to answer.

For example, if we want to train an ML model to predict customer churn, we can feed the model with historical data about our customers, including age, gender, geographical location, purchasing behavior, and product usage. Crucially, we augment this dataset with the churn data for each customer. Churn, in this context, refers to when a customer stops using a product or service. Once trained, the model will be able to forecast the date when a particular customer is likely to churn based on the available data for that unique customer.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

- Reinforcement Learning (RL) addresses a key limitation of supervised and unsupervised learning: their inability to discover novel mechanisms beyond the relationships in the data. Reinforcement Learning algorithms such as deep-Q learning, or actor-critic, enhance effectiveness by allowing the exploration of strategies not represented in the available data. This approach is akin to “A/B testing on steroids,” where millions of different scenarios are automatically generated, and the feedback loop is integral to the learning process.

For example, the machine might experiment with varying discount levels to identify the most effective discount at each stage of the customer journey. The RL system might issue random discounts to thousands of customers and then measure the churn rate. This process refines the discount strategy to optimal levels based on the knowledge of “what works or not” for specific customer groups. For example, the system might learn that customers with children value a free month in August, while affluent customers in rural areas respond well to bundle discounts. In this way, reinforcement learning enables the discovery of effective, tailored strategies that were not initially apparent in the data.

- Unsupervised Learning refers to an ML approach where models are trained on unlabeled data. Unlike supervised learning, where we have explicit labels for our data, unsupervised learning involves discovering underlying structures without predefined categories. In this scenario, there aren’t specific targets to predict, but we aim to uncover meaningful patterns, similarities, or clusters in the data.

For example, by feeding customer data into the ML solution, we might discover that customers without children exhibit similar behavior across different geographical locations and income levels. Armed with this insight, specific marketing campaigns can be developed to target this particular customer segment effectively.

AI's limitations

The current limitations of AI systems include hallucinations, bias, explainability, and scalability:

- **Hallucinations:** AI hallucinations refer to inaccurate or misleading results generated by AI models. These errors can arise due to various factors, including limited training data, faulty assumptions made by the model, or biases present in the training data.
- **Bias:** AI systems are trained on large datasets, which are created by humans. Consequently, these systems can inadvertently learn and perpetuate the biases present in the training data. Common biases include:
 - **Cognitive biases:** Unconscious errors in human judgment that can influence AI training and outcomes.
 - **Implicit bias:** Unconscious stereotypes that can be reflected in AI if present in training data.
 - **Gender, racial, and age bias:** AI decisions can unfairly disadvantage people based on these factors.
 - **Sampling bias:** Occurs when training data isn't representative of the population.
 - **Temporal bias:** Occurs when training data is time-specific and may not reflect current or future states.
 - **Over-fitting to training data:** Occurs when AI performs well on training data but poorly on new data. This happens when a model captures external factors (noise) that alter the data points but are not part of the system that we want to model along with the underlying pattern in the training data, resulting in poor performance on new data.
 - **Algorithmic bias:** Prejudiced assumptions made during algorithm development can lead to biased outcomes.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

- **Explainability:** AI models, including LLMs, are often referred to as “black boxes” because it’s challenging to understand how they arrive at a particular output. This lack of transparency applies in particular to neural networks and deep learning systems whereas simple ML algorithms like decision trees allow an external user to understand how the AI came to a specific result.
- **Scalability:** AI systems that perform at scale in production settings can become very expensive and very hard to maintain for a number of related reasons:
 - **Data sourcing and ingesting:** In general, large AI models perform better when more good quality data is used to train them. Large, good quality datasets are difficult and expensive to handle and store.
 - **Compute cost to train:** Most deep learning systems, and LLMs in particular, require vast amounts of compute power to train, with periodic retraining to maintain acceptable performances or to add more recent data.
 - **Compute cost to operate:** Running a model that serves many concurrent users can become complex and expensive in exponential ways.

Copyright and AI

A complex, critical issue linked to current limitations of AI is that of copyright infringement. As I write this book, in 2024, it is still unclear how the issues of consent, credit, and compensation can be addressed and resolved. Most authors argue that before an LLM is trained using their original work, they should be consulted to give permission. Then, when their work contributes to an LLM output, they should be credited and compensated.

Most AI companies argue that it’s impossible to correctly attribute credit since LLMs are not explainable and it’s impossible to determine what data source is used to generate arbitrary output. They also argue that LLMs and generative AI in general are not dissimilar to artists who draw inspiration and build on the work of those before them. It’s difficult to predict where the copyright law is going to settle, but recently large publishing companies and AI organizations have signed contracts to regulate the usage of copyright-protected material. It

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

seems likely that more of those deals will be struck to try and create a viable business model.

The issue of copyright is gaining momentum. However, there are conflicting and contradictory regulations across different regions. As an area of ongoing change and development, business leaders are advised to refer to the latest regulations in the regions where their organizations operate.

Trustworthy AI

Trustworthy AI is also an emerging field originating from collaborative international efforts, privacy-enhancing technologies, and ethical principles, aimed at reliable, transparent, and privacy-respecting AI systems. It has become an integral part of the AI system lifecycle to ensure beneficial and reliable outcomes. It requires governance and regulatory compliance, from ideation to design, development, deployment, and machine learning operations.

Generative AI and LLMs in more detail

LLMs, a subset of generative AI, are good at summarizing context and language, which is particularly effective in text-based writing assistance, content creation, and reasoning over structured and unstructured data. Companies are applying LLMs across different sectors to enhance their operations and we have in-depth case studies where we discuss these applications.

The common emerging applications for generative AI include in operations, enabling a new type of customer engagement and care, empowering advisors and market researchers, building co-pilots to help customers find knowledge faster, improving marketing effectiveness, and accelerating product development.

Foundational components of LLMs

Tokens and vectors, embeddings, and transformers are all interconnected concepts and play crucial roles in processing and understanding natural language data. They help LLMs to grasp the intricacies of human language

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

and perform a wide range of Natural Language Processing (NLP) tasks. Here's a brief explanation of each.

Tokens and Vectors

In the first stage of the process of understanding human language, the text data is broken down into smaller units called tokens, which can be words, parts of words, or punctuation marks. These tokens are then transformed into vectors, which are numerical representations of these words. This transformation process allows the model to understand the meaning of words and how they relate to each other.

Embeddings

The vectors created from tokens are also known as embeddings, which make text, images, and audio digestible for AI transformers. Embeddings are numerical vectors that represent, in a machine-readable format, the semantic context of a text (for chatbots) in a format that is language-independent and designed to retain as much knowledge as possible. For example, the machine will be able to parse the token "Florence" as a city (the Italian city in Tuscany) or a first name (as in "Florence Nightingale"). This allows the machine to retain the implied knowledge that in a specific text might be describing the architecture of a city as opposed to the founder of modern nursing.

A key feature of embeddings is that similar concepts have similar values. This feature enables search, comparison, and manipulation beyond simple full-text queries, thereby enhancing the capabilities of AI systems.

Transformers

The transformer is the core engine of LLMs. Its primary function is to convert an input embedding into an output embedding, which is then "translated" into human language. The transformer leverages all its training to determine the most probable output given the user's input. It generates its output by selecting the most likely token at each step. This process allows the transformer to produce

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

coherent and contextually appropriate responses. It's a key component in the effectiveness of LLMs in understanding and generating human-like text.

The evolution and impact of transformers in AI

In 2017, a paper titled “Attention Is All You Need” was published by Google and introduced the concept of transformers. This kick-started an entirely new era of AI. Transformers now form the backbone of many cutting-edge AI applications enabling the scaling of LLMs like those that power ChatGPT and Gemini.

While LLMs focus primarily on language, generative AI models can serve other purposes, transforming data between multiple modalities, including text, images, video, and audio. Regardless of the purpose, each model interprets the underlying meaning of its input and generates the most likely output according to the training data.

The real power of LLMs comes from the fact that they work in areas far beyond language. The latest release of AI solutions is multimodal, meaning that they can use a text prompt to generate an image or video. Examples of multimodal systems include:

- DALL-E, an OpenAI system that can create images from text descriptions, using a large neural network trained on a dataset of text-image pairs.
- Midjourney, a generative artificial intelligence program and service that can create images from text descriptions (prompts).
- Stable Diffusion AI, an open-source AI model that can generate realistic images from text descriptions or other inputs. It uses a technique called latent diffusion, which involves the iterative addition of “noise” – random variations or distortions – to an image until it matches the desired output.
- SORA OpenAI system generates videos from text prompts, due for launch in 2024.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

Model Customization – refining existing models

The current LLMs available on the market, like GPT, are pre-trained on a diverse range of internet text and can be used by many users at the same time. However, they are not designed to learn or adapt to new information after training. This means that while they can generate creative text based on the input provided, they can't be fine-tuned on a specific task without going through a re-training process.

For instance, if you have a specific set of documents on a specialized topic and want to teach an LLM to understand and generate text about this topic, you cannot do this with models like ChatGPT. Instead, you would need to take a general-purpose LLM and train it on your specific dataset. This process, known as fine-tuning, adapts the model to your specific needs. However, fine-tuning a model is technologically complex and expensive. The investment is typically justified by the potential return or by addressing concerns such as security and privacy that require the model to have a deep understanding of a particular domain. A practical example of a fine-tuned LLM is a chatbot trained on the FAQ of a large enterprise. This chatbot can answer specific questions about the company's products or processes, providing valuable assistance to both customers and employees.

Grounded Generation/Retrieval Augmented Generation (RAG)

Grounded Generation, or Retrieval Augmented Generation (RAG), is a secure, cost-effective alternative to fine-tuning in LLMs. It reduces hallucination levels and uses semantic retrieval techniques for context provision. It can provide citations, enhancing the credibility of its outputs. RAG is adept at extracting, encoding, and indexing intelligence from various sources, which is then retrieved and ranked per user queries. It generates relevant content by considering unique contexts and historical materials. The choice between Grounded Generation and fine-tuning depends on the task requirements. Some projects may benefit from a combination of RAG and fine-tuning, enhancing model performance and reliability. This trade-off is explained in depth in the Vectara case study (*Chapter 26*).

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

The commercial opportunities of the democratization of LLMs

Moving forward, as demonstrated by the companies in this book, the expectation is that businesses will take a novel approach to solving new problems previously labeled unsolvable. Today, there are many people who until now didn't have access to this capability nor come from AI research domains, but who now have access and can bring their creativity to solving new problems.

Multiple companies are competing to establish and develop environments that can effectively work with open-source AI models, including GPT, Claude, Gemini, and Llama among others. This open-source nature is leading to diverse and innovative applications of these models, as individuals and organizations take them in unique and interesting directions, as we will see in the following interviews and case studies.

The pace of development in this area is likely to be rapid, with companies striving to create use cases that offer some form of differentiation. This competitive landscape is expected to drive investment, innovation, and progress in the field.

How controlled your industry domain is will heavily influence how applicable the models will be. In areas of very high regulation, such as in finance and healthcare, where being correct versus close enough is most important, these models are less applicable without considerable investment in customization.

2.b AI Technology Stack

An AI technology structure, often referred to as “the stack,” is a set of tools and technologies that are used to develop, deploy, and manage AI applications. It consists of different layers that perform various functions such as data collection, data processing, machine learning, model training, model deployment, and user interface. An AI tech stack can vary depending on the type and complexity of the AI project.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

To understand the components of the tech stack, let's use an example of an infrastructure investment firm wanting to forecast the value of their assets at exit. To accomplish this, the firm will need the following:

- **Data sources:** Both internal (for example, asset depreciation history) and external (for example, macroeconomic factors, demographic data) sources to provide the necessary data.
- **Data pipeline layer:** Automated processes to ingest data from various sources and store it consistently in a central data store.
- **Data storage:** A system to efficiently store and potentially transform the data to meet enterprise needs.
- **Data quality:** Ensure datasets are of good enough quality to properly train the AI model.
- **AI algorithm:** Use the data from the storage to train a forecasting model to predict asset value at exit. Past asset valuations and market conditions contribute to the training of the algorithm.
- **AI platform:** Provide a toolset to efficiently build, test, and deploy AI algorithms in an enterprise, shared environment. The toolset includes Integrated Development Environments (IDEs), code editors, debuggers, testing tools, version control systems, and documentation tools.
- **APIs (application programming interface):** Forecast data is accessible programmatically through an internal system. APIs allow authorized third-party systems to access the data in a controlled and monitored way.
- **AI solution:** A fully integrated solution that allows business users to explore and analyze forecasted asset valuations as well as “what if” scenarios to anticipate various outcomes. It often includes a web-based interface for interaction and a dashboard for standardized reporting.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

Data as a universal resource

Without data, there is no AI. The emerging imperative for businesses is to not only adopt AI models but to also establish robust data management practices. Historically, some companies may have managed without a proper data architecture, but to leverage AI effectively, a solid foundation is essential.

Despite the buzz around AI and machine-learning techniques, their applications are still maturing, with the majority of effort and time being spent on the data foundations which are needed to enable data usage. We will encounter this challenge in honest interviews with practitioners, particularly in more established companies.

Many companies seldom question the value of their data. However, when companies invest in structuring their data and using it effectively, there is a significant shift. Proper data management can streamline operations, improve decision-making, and ultimately drive growth.

Data comes in many forms, all of which are valuable for AI systems. The following are examples of data types that often serve initial use cases:

- Transactional data includes data about sales, purchases, and other customer transactions. It's often used for forecasting and personalization.
- Financial data includes sales figures, revenue, costs, and other financial metrics. It's often used for forecasting and financial planning.
- Customer data includes demographic information, purchase history, and other information about customer interactions. It's often used for personalization and churn prediction.
- Customer support data is generated from customer interactions with support services, including call logs, emails, and chat transcripts. It's used to improve customer service, identify common issues, and enhance product/service quality.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

- Operational data includes data about business operations such as supply chain data, inventory data, and other operational metrics such as inventory, suppliers, manufacturing process and efficiency, shipments (timelines and costs), procurement and compliance. It's often used for recasting and optimization.
- Web analytics include data about website usage, such as page views, click-through rates, and time spent on site. It's often used for personalization and understanding customer behavior.

In addition to the typical data inputs, several other types of data are widely used to support various use cases such as personalizing customer journeys, B2B acquisition, understanding customer sentiment to inform marketing strategies, optimizing supply chains, and new product development. These include:

- Sociodemographic data, which includes information about individuals like age, gender, and education. It's essential for market research, customer segmentation, and tailoring products/services to specific demographic groups.
- User journey data, which is generated by tracking the path customers take through a digital channel (website, application, chatbot or conversational interface). It includes data points like page views, clicks, page browsing times, questions asked, sentiment (answer was useful versus not), basket adds, and purchases.
- Firmographic data relates to organizational characteristics and typically includes specifics such as industry classification, company size gauged by revenue and employees, geographical locations, investment profiles, fundraising history, credit ratings, and so on. It's used for market segmentation, targeted marketing, and competitive analyses.
- Social media data, which includes user-generated content on B2B and B2C social media platforms, including posts, likes, shares, comments, and engagement with content.

OVERVIEW OF AI CONCEPTS AND TECHNOLOGY STACK

- Product telemetry data, which is collected from products or devices, providing insights into usage patterns, performance issues, and user interactions. This data is typically used for product development, quality assurance, and customer experience improvement.

Data readiness

Data needs to be prepared or pre-processed to be suitable for the specific AI model. Data readiness refers to the extent to which a business's data is prepared, available, and suitable for use in AI applications. However, achieving data readiness is often challenging due to data assets scattered across different departments and systems, poor quality, lack of standardization, and accessibility as well as gaps in processes for compliance and security measures.

Lack of data readiness is often the primary factor contributing to delays or abandonment of AI initiatives. Therefore, ensuring data readiness is crucial for successful AI implementation.

AI as a Service (AIaaS)

Building and then maintaining AI solutions can be costly and time-consuming to satisfy compliance standards and regulations. AIaaS enables customers to access and use cloud-based AI capabilities without having to invest in the hardware or software needed to develop and deploy AI applications. It provides scalability, allowing clients to adjust their resources according to their needs without worrying about infrastructural limitations. AIaaS ensures flexibility as customers aren't required to have expertise in AI or machine learning and can select from an array of pre-built or tailored solutions. Furthermore, it enables access to the latest technologies and best practices in the field.



3

Sam Liang

CEO of Otter.ai

We want to build a large spoken language model that can generate realistic conversations. Eventually the OtterPilot will have a really natural conversation with one person, or even with multiple people, and really understand the dynamics at play. In a few years, it's conceivable that OtterPilot will actively participate in meetings... Quite fascinating, isn't it?

LinkedIn: [samliang](#)

SAM LIANG

Sam Liang, CEO of Otter.ai, followed his initial studies in computer science with a PhD at Stanford University in computer-distributed systems. Instrumental in the development of the Google Maps location platform, Sam engineered the back-end system that powers Google Mobile Maps, used by billions of users worldwide. In 2010, he ventured into the world of entrepreneurship, dedicating his time and knowledge to building startups.

His innovative next venture in Palo Alto, Silicon Valley, initially focused on the development of a mobile analytics system. This caught the attention of Alibaba, leading to an acquisition in 2016. In that same year, undeterred by the fact that conversational AI was in a nascent form, he launched Otter.ai – a speech recognition system that is now a transcription service used worldwide.

LISA: How does your AI-powered system enhance business collaboration?

SAM: Our AI has high accuracy and speed, operating in real time, recording and processing information as it happens. The end product is live-streamed, providing users with immediate access to the data. The AI's output appears on your computer almost instantly. But that's not all. In the background, it's also comprehending the context and understanding the nuances of the conversation. Meanwhile, our AI is also generating a real-time summary.

We have been working with AI for seven years now. During this time, we've been diligently building an AI-powered collaboration system for enterprises. Working in a large enterprise requires you to spend a ton of time in meetings – people forget how much time they are talking and listening on a daily basis. So, a tremendous amount of investment, both in terms of time and resources, is allocated to meetings. We're trying to disrupt the current system, to optimize meetings, and reduce unnecessary ones. Using Otter, many large meetings can be bypassed as it swiftly provides a content summary, eliminating the need to sit through 60-minute meetings. This is a whole new world. We're really excited about ChatGPT and other large language models and we're building our own AI to enhance the value of conversations.

LISA: What was the motive behind the origin of Otter.ai? Could you share the story behind its inception?

SAM: There were two motivations. The first is my deep interest in data as a computer scientist. Think about the fact that humans have been around for about 100,000 years, and for most of that time, we've communicated verbally. But no voice data was saved, so that's a huge loss of human knowledge and intelligence. It was only after Thomas Edison invented the audio recorder that we began to capture voice, but even then, 99.9% of human conversation is never recorded. So, I realized that there's a big loss of knowledge there.

On a practical level, as a startup founder working with businesses and having to go to so many meetings, I found it hard to remember who said what or even recall what I had promised others, and this was back in 2016. I found it frustrating that I could use Gmail to search for anything from 20 years ago, but I couldn't search for something I had heard just three hours ago. This led me to think, why not build a product application that enables this?

So, I had two motivations. One was about wanting to collect as much voice data as possible. The second is the practical need to help people remember the information they discussed and help teams share the information relevant to a broader audience outside the meeting. So, how do we share this information effectively? We wanted to solve that problem.

LISA: What are you focused on now to build out the services?

SAM: We're building a horizontal product to help people with their meetings across different domains. This includes investment meetings, board meetings, internal team meetings, and even customer support or sales meetings.

We've recently launched a new product called OtterPilot for Sales. It's a vertical solution built on top of our horizontal platform. This solution offers specific features tailored to the sales use case, for when a sales rep speaks to the customer.

SAM LIANG

In a typical sales interaction, the sales rep needs to listen very carefully to the customer's requirements and pain points, the competitors they're considering, their budget, and pricing requirements. Traditionally, the sales rep would need to carefully take notes and then, after the meeting, manually enter this information into a CRM system. Now, with AI OtterPilot for Sales, we completely automate this process because AI can listen, with the participants' permission, transcribe and save everything, and then identify the important insights that need to be saved into the CRM. The AI extracts this information and instantly enters it into the CRM, saving the sales rep valuable time.

The AI provides transparency. The sales call information is available to the sales manager as well. The sales manager's job is to monitor how all the sales calls go and to keep coaching the sales rep. Traditionally, the sales managers actually don't have visibility on how all the sales calls go but with Otter, they can access every call, all the recordings, transcripts, summaries, and even get a score for the call. This automation and insight generation can really help both junior sales reps and sales managers, potentially uncovering insights they might have missed.

LISA: Can you bring this to life with a customer case study?

SAM: Certainly. We worked with a US-based IT automation services company facing productivity challenges due to the cumbersome process of manual note taking during sales calls. Otter helped them to:

- Streamline the note-taking process to allow reps to focus on customer interactions.
- Automatically record, transcribe, and store meeting notes to improve efficiency.
- Easily tag and assign action items to participants to increase collaboration and move deals quickly.

SAM LIANG

The impact was significant: Otter greatly improved workflow, saving the sales team approximately one-third of their time and increasing productivity.

Another use case was the work we did with a global financial media company to address various challenges, including those arising from the increase in remote work and the need to ensure inclusivity within their workforce. The media company decided Otter was the best company-wide solution for its tens of thousands of employees due to its:

- Speed and accuracy over other AI-powered transcription tools.
- Searchability feature, which allows media journalists to easily find what they're looking for.
- Ability to save time in meetings, give journalists confidence in their notes, and provide accountability.
- Potential to foster a more inclusive workplace, as Otter offers employees with disabilities new options for accessing their work and communicating in the workplace. For instance, Otter offers employees with hearing impairments an essential solution.

LISA: Could you elaborate on the technology stack behind Otter?

SAM: Our technology is designed to understand and transcribe conversations with high accuracy and speed in real time. Most of the products you see today don't provide live transcripts, or handle accents, names, terminologies, or medical terms if you're in the healthcare domain. These are all challenging and a lot of hard work for AI to do. We've also developed technology to understand the voice print of each person and identify who's talking. It's natural for humans, but challenging for machines, to separate one person's voice from another.

The raw transcript could have thousands of words from one meeting, so to make it more consumable we summarize and condense the information into a few bullet points and identify the key topics. You can then search for topics and find a few bullet points for each one.

Moving forward, we're building a large spoken language model, which is different from the models built by OpenAI and others. Our focus is on training with conversational data, as spoken language and human conversations have different characteristics compared to written documents.

The most prominent differences involve multiple speakers who, as they take turns, will each speak with specific intonation and emotion. This dynamic is special to conversations and differs from written documents, which are usually formal and have a well-designed structure. Conversations, on the other hand, are free-flowing and unstructured, and the presence of multiple speakers makes things even more complicated.

The next step is to actually generate conversations. We want to build a large spoken language model that can generate realistic conversations. Eventually the OtterPilot can have a really natural conversation with one person, or even with multiple people, and really understand the dynamics at play. In a few years, it's conceivable that OtterPilot will actively participate in meetings, using advanced AI analysis to identify key topics, actions, decisions made, and discussion sentiment. Additionally, it could provide multilingual support. Quite fascinating, isn't it?

LISA: Is OtterPilot being developed only in English or in multiple languages?

SAM: At present, we are focusing solely on English. However, we have plans to extend our reach to all other major languages in the future.

LISA: How does Otter's service differ from other team collaboration platforms, particularly in terms of transcription features?

SAM: Microsoft Teams offer transcription services, and has recently announced a co-pilot, so there is some overlap, but our high-level vision is different. Microsoft Teams has been around for a while, starting as a video conferencing system with added messaging capabilities. But when it comes to the voice component, even though they provide transcripts, our goal is different.

We aim to build a system that provides transparency for the entire organization. The sharing model is really important. Our philosophy is that at most meetings, competencies should be open and transparent, regardless of whether you're part of the sales, marketing, engineering, or design team. We estimate that about 90% of meetings should be public, to facilitate information sharing and flow, and to reduce information silos, which often lead to inefficiencies within companies.

Even within our own company, we see that sometimes one person discusses the same topic with three different groups of people. To reduce this inefficiency, perhaps it would be more effective to have the discussion once and then share that with the other two groups.

We're in the process of building a collaborative AI-powered system, and meeting transcription is just one small part of it. As we grow and evolve, others may start to emulate us, just as they did with the business messaging app, Slack.

While Microsoft certainly controls a significant portion of the market, you always see startups disrupting larger players. Zoom is a prime example. Even just five years ago, it was hard to imagine Zoom capturing such a large market share, especially considering that Microsoft Teams, Google Hangouts, and Webex had been around for quite some time.

SAM LIANG

LISA: How would you characterize your customers?

SAM: Our customer base is relatively horizontal, composed mainly of knowledge workers. We operate with a PLG (Product-Led Growth) approach, which means most people sign up independently.

LISA: Given the broad application of your product, are there specific knowledge professionals across various industries and company sizes that particularly benefit from it?

SAM: Our user base is pretty diverse. I will say we've seen knowledge workers from all kinds of industries using Otter – they could be salespeople, consultants, writers, and more. We have a mobile app, which is a key part of our strategy and offers a different experience from Microsoft Teams or Zoom. Users can record their thoughts while walking in the park or even driving. Writers often use Otter to capture their thoughts. Then, Otter's edit mode allows users to correct errors and adjust line breaks for cleaner transcriptions. Finally, they can copy and paste the text into the articles they're writing. So, we see a wide range of use cases.

LISA: How have you addressed concerns about security, particularly given that your product is used in people's meetings?

SAM: Suspicion or fear is understandable. The idea of being recorded is a new thing. In the past, it usually indicated something bad, like being in a police station. You're talking to the police and they might record you. Times are changing, and our perceptions need to adapt. It's a balancing act.

Think about it – do you want AI to assist you in remembering details and sharing information with your team? AI can be a great help, but it needs data to work with. Of course, this means you have to keep it confidential and secure. There is no question about the security model. We didn't invent anything new – this is pretty standard across the industry in terms of cloud-based services – Microsoft Teams, for instance, does the same thing. If you record a meeting

or generate a transcript, it's stored in the cloud. Zoom operates on the same principle. Even when you use Gmail, your emails, confidential as they may be, are stored in someone else's cloud system.

Otter is similar. We're no different from those other services. Of course, we need to overcome some conventional nervousness to use such a new tool, but things will change.

Key Takeaways

- **Strategic advantages of LLMs:** Investing in LLMs can lead to innovative solutions that improve business operations, including augmenting tasks and generating new perspectives. Those firms that have been investing in AI for multiple years have a head start with generative AI. It requires time to build muscle and incorporate AI into work processes.
- **Benefits of AI-powered transcription services:** AI's high speed and accuracy provide improved efficiency and transparency by recording and processing information as it happens, as well as generating key insights. Using advanced AI analysis, these solutions are moving toward more active participation in meetings and creative collaborations. For example, pioneering artist will.i.am is actively experimenting with a new AI cohost for his podcast.
- **Reducing inefficiencies in collaboration:** As companies scale, particularly those that employ a significant number of knowledge workers, inefficiencies tend to increase due to the proliferation of time-consuming meetings and information silos. LLMs, particularly AI-powered transcription and audio tools, can play a crucial role in addressing this issue.
- **AI solutions for enhanced inclusivity:** AI solutions have proven to significantly broaden accessibility for employees with disabilities. They offer new opportunities for effective communication and work access, which can foster a more inclusive work environment.

- **Product-Led Growth (PLG) strategy:** This relies on the product as the primary way to acquire, engage, and retain customers. Product adoption scales through user experience rather than a sales force. The strategy encourages individuals to sign up independently, fostering organic growth by word of mouth and referrals. It highlights the importance of an intuitive product that can be used without extensive guidance or training. This approach is instrumental to the product development of AI solutions that require high user engagement and data-driven iteration to evolve the product.



4

Amr Awadallah

Founder and CEO at Vectara

Generative AI is transformative because, for the first time in history, we have AI systems that can interpret human language as we do, which is unprecedented.... Minimizing hallucination, bias, and copyright abuse are key issues that anybody thinking about deploying this technology cares about.

LinkedIn: awadallah

AMR AWADALLAH

Amr Awadallah is the Founder and CEO of Vectara, a generative AI platform that offers a complete, modular, and API-based search pipeline for developers. Originally from Egypt, Amr moved to the US in 1995 to pursue a PhD at Stanford University after which he launched his first startup which Yahoo acquired. Following an eight-year tenure at Yahoo, Amr launched Cloudera, a pioneering big data company that went public in 2017 and was subsequently acquired by a private equity firm. Amr subsequently served as Vice President of Developer Relations at Google Cloud for two years before starting Vectara, a venture dedicated to building tools that enable enterprises to leverage generative AI technologies, using their own data.

LISA: What is Vectara's core product and how does it enhance the capabilities of generative AI technologies for business?

AMR: Vectara is a pioneering AI platform that empowers organizations to harness the power of generative AI using their proprietary data to innovate and scale. This integration ensures that interactions with the large language model (LLM) are uniquely tailored, providing responses that reflect the company's own information and context. In addition, Vectara addresses common challenges associated with AI models, such as reducing hallucinations – the generation of inaccurate or nonsensical information – mitigating bias and preventing copywriting violations. This focused approach makes Vectara an invaluable tool for companies looking to leverage AI while maintaining accuracy, relevance and legal compliance.

LISA: How has the software industry been driven by efficiency and cost reduction, and how have historical applications of AI contributed to this advancement?

AMR: Software is all about efficiency, whether that's in our business and our processes, our finances, or our marketing. How can we do it with less cost and with better performance, meaning less human error?

The AI evolution that we're going through right now – generative AI – is a subset of the bigger AI evolution that has been going on for some time. In the previous AI wave, there were three primary subsets. The first was pattern recognition AI applications, which have been pivotal in detecting fraud, identifying quality issues in products or pipelines, and diagnosing diseases. Pattern-based AI is a facet of machine learning that focuses on identifying previous patterns and regularities in data that indicate for example, this type of disease or this type of failure or this type of fraud, and then we just want to find it again at scale.

The second subset was clustering algorithms that tailor user experiences for more effective marketing - the personalization example. Can we cluster the behaviors of groups of products, groups of users, and so on? And can we either upsell them, cross-sell them, or market to them in the right way with the right message?

The third subset was forecasting future scenarios through time series analysis. For instance, given the history of this sensor on this boiler, can I predict the next time this boiler is going to fail? Or, given this stock, the high-frequency trading stock price, how it's moving and the signals I have, can I predict what that is going to look like in the next few minutes? These were the primary three subsets in the previous AI wave.

LISA: How is generative AI different from the previous AI wave, and what benefits do you see it bringing?

AMR: Generative AI is transformative because, for the first time in history, we have AI systems that can interpret human language as we do, which is unprecedented. It represents a milestone in the user interface between humans and software.

The analogy I share about this moment is the leap from command-line interfaces to Windows 95 and then the iPhone. Originally, the only way we could interact with software was through something called the command line. If you are old enough to remember DOS, then you remember all the commands and options you had to put after the commands. It's a fixed set of words and you

have to use them in that way. Very few people knew how to do that right, so the consumption of software was limited to very few people.

Then Windows 95 introduced the concept of menus and a mouse that you can move around. So, instead of remembering the commands, you can pick the command on the menu. The forms were a lot more intuitive with nice user interface modalities, which open up the consumption of software for consumers or inside of businesses, or to a bigger class of people. However, we literally had to rewrite all of our software to adapt every single application that worked in the command line to work in this new way. That was the first phase of how we evolved our interaction with software.

The following phase was the iPhone and touch screens, which opened up the consumption of technology to many more people because everybody knows how to use their fingers. However, it also meant, again, that we had to rewrite every single application that we used both in and outside of the enterprise. We had to evolve it to leverage this new modality of interaction.

Today, generative AI technologies allow us to interact with software in the same way we interact with each other. We can introduce the software to this particular task or assignment, and it can ask “What do you want me to help you with today?” You can just say it verbally or write it down, and the software will figure out how to do it for you. Of course, it’s not perfect yet, but we have the beginnings of that. The consumption will become ubiquitous; everybody will be able to use software to do their tasks. It will be just like in the previous waves, and that’s the trend trajectory my company is betting on. Every single software application in the next 5 or 10 years will need to be rewritten to leverage this new modality of interaction.

That’s what we are about enabling and asking, “What are the benefits of this wave?” Of course, it’s going to depend on the use case and the application, but the common benefit all of these things will have is efficiency. Whatever used to take us hours or days to do, we’ll now be able to do in minutes.

LISA: How does generative AI transform the search for information?

AMR: We are moving from the old method which was about search engines where you search documentation and the web, to get back a list of results or a bunch of videos you can watch to get something done. Now, instead of a list of results, we can just have the answer.

For example, I might ask: “How do I take a vacation using Workday?” And it would say “Step number one, you do this, step number two, you do that” and so on until you’re done. If we get that answer correct every single time without hallucinations, then we can move from answer engines to action engines, which means it would say, “Step one, do this, step two, do that, and do you want me to do that for you?” And if I say yes then it reasons with you to understand the context and guides you. It completes the task without you having to interact with any forms, without you having to learn how the software works. It adapts to you and your style and how you’re answering the questions to get the task done. And that’s the journey we’re on. We’re moving from legacy, which is where we had to search for information to find how we can do something, to answers. It’s a move from having the answer for “How can we do it right away?” to where it just does the action for me so I can move on.

The action can be a simple task like taking a vacation using Workday, or the action could be “I’m looking at these two lawsuits right now. I want a legal draft of how this lawsuit is going to be structured. Look at all the previous lawsuits that were similar to those I wrote, or ones that other lawyers and other teams wrote, and compare that to the judges and what they ruled, how they said yes or no, and whether it succeeded or not. Then summarize all of that into a new legal draft so that I, as a lawyer, can start working now to finish my task.” The lawyer now finishes doing the draft in literally 15 minutes versus 5 weeks of research before. This is really what we’re talking about – efficiency. Everything is going to be 100 times more efficient.

Those of us who know how to embrace that efficiency in our jobs are probably using ChatGPT in one shape or form right now. But those of us who know how to embrace it and leverage it will be 100 times more productive in everything we do, and those of us that don't will fall behind.

LISA: Which tasks does AI now perform better than humans?

AMR: The fear of job loss is not because of AI but, as I like to say, because of those who know how to use AI versus those who don't. Ultimately, job security will depend on the ability to leverage AI although there are some roles, like customer support and driving, which may become obsolete due to automation. Five years from now, I foresee there won't be humans working in call centers or customer support. And in 10 years, there will be fewer human drivers. Today, in San Francisco you can call a car with an app that will arrive without a driver, pick you up and drop you off where you want. In my opinion, it's the perfect driver, and represents a significant advancement in technology. I envisage that tasks such as these will be completely replaced, but in the majority of cases, it's going to be people who know how to leverage the technology to their advantage who will keep their jobs.

LISA: How will the leading technology companies impact generative AI, and what's your differentiation strategy?

AMR: Everybody is waking up to the power that generative AI offers. The big players are all going to have eyes on this. When I was building Cloudera to do big data analytics it was the same thing. Big data analytics was a big field, and the big players were all there. It did not prevent me from growing Cloudera to be a \$5 billion company. It did not prevent Snowflake from building a \$50 billion company, despite the big players having some of the best databases; for instance, Microsoft has SQL, Cosmos and Cosmos DB, and also provides Azure Active Directory for identity management. AWS has Redshift and Google has BigQuery offering data warehousing solutions. Despite the big players having all these, Snowflake is still able to succeed. In other words, this is a huge, huge

marketplace and tons and tons of investment will happen over the next few years. I'm not worried by the big players jumping into this same space.

For many companies in this world, IT is not their thing. They are about manufacturing, they are about healthcare, they are about law and so on. Because they're working on contracts, they don't have the engineering teams to go and build something as powerful as this. So, Vectara's key differentiation is about ease of use. How can we build a system that any IT person managing the laptops of a company can integrate into their workflows and start leveraging it to do something? We offer a very, very simple, easy-to-use platform from end to end.

It is similar to Snowflake. They came in with simplicity and it worked very well for them. But we also have some core IP around the hallucination issue that I mentioned. Minimizing hallucination, bias and copyright abuse are key issues that anybody thinking about deploying this technology cares about.

LISA: How does your platform provide a unique advantage to its users?

AMR: We integrate a ChatGPT kind of experience by making it a function of the proprietary data you have in your organization. For instance, if you're a lawyer with an extensive archive of past contracts, our platform enables you to utilize that wealth of information in current legal work. Similarly, a PR firm can leverage a history of press releases, and a bridge construction company can access a repository of construction plans, workflows, and supply chain details. Even an investor can input a multitude of investment pitches and memos into the system. Then, when you ask it a question or set a task, it's answering not because of its knowledge of the world like ChatGPT does, but rather because of the knowledge of what you have.

So, the key strength of our platform lies in its ability to provide context-specific responses informed by your organization's accumulated knowledge. And because its responses are grounded in your data, it doesn't have to hallucinate as much. Hallucination arises because it doesn't know what it doesn't know but, by always grounding it in your own data, it reduces hallucination.

LISA: If it is only using the data from that company, how can it be as smart and responsive as ChatGPT?

AMR: That's a very good question. The process is analogous to human learning. Consider a college education where you absorb a vast array of information from books and lectures. Similarly, we train large language models on extensive datasets, distilling the essence of that information into a compressed form. This is akin to summarizing key concepts in your memory without retaining every detail. When recalling information, sometimes your brain fills in the gaps with what it believes to be correct, which is similar to what we refer to as hallucination in AI. Despite compressing internet-scale data by a significant factor, the resulting model can still respond intelligently within its knowledge scope.

LISA: When compressing data, how does the model determine what information is valuable and what's not?

AMR: Filtering the data involves two primary methods. Because the quality of the source data significantly influences the model's performance, a human curator will initially review, refine, and clean up the input data to guide the model. Post-training, we employ a technique called Reinforcement Learning from Human Feedback (RLHF), where the model learns from its errors as humans correct its misguided outputs, saying "Don't say this again, bad model." While not infallible, this method helps the model avoid rubbish or irrelevant areas in its "brain," reducing inaccuracies.

LISA: How does the platform ensure accuracy when working with a company's proprietary data?

AMR: The solution is that you have to retain the original data rather than compressing it, akin to a model with a photographic memory. The AI has a foundational understanding of language (for example, English, German, Chinese) and domain-specific concepts (legal, manufacturing, engineering,

health, and so on) so it has the framework to read, interpret, reason, and process the information but it doesn't have the information itself. It relies on the factual data provided by the company for generating responses. When queried, the AI correlates the question with the relevant facts before generating a response.

In this way, we are grounding the solution in facts that come directly from the proprietary data of that customer or user. This significantly minimizes hallucination because we explicitly tell the model to think about it in a very specific way, for example, with a prompt like, "You are a financial analyst. Only answer this question. And the question is, should I invest in this company or not as a function of these facts?" And then the facts are the previous decisions similar to this decision.

We would also say: "If the facts do not help explain this question, say 'I don't know,' don't make something up. Say literally 'I don't know how to answer the question.' If the facts are relevant to the question, then please summarize them together in a cohesive response and cite every single fact that you used to come up with a conclusion in the response so we can get the confidence you're not making stuff up." It's like the fact-checking that newspapers do with articles.

LISA: Given the skill and expertise of creating accurate prompts, are you tailoring your product to specific vertical use cases for more effective use?

AMR: I have two answers to that. The first, the visionary perspective, takes a page from history. When breakthroughs like the iPhone and Windows 95 came out, predicting the full spectrum of apps they'd enable was beyond our reach. For instance, we anticipated that many of our basic apps had to transition to the iPhone to work with a touch screen, but innovations like Snapchat or Tinder's unique swipe right and left were unforeseen. So, with that caveat, I acknowledge there will be unforeseen applications of our product as well.

On the practical side, we're gaining traction with certain applications. For instance, every business needs one-on-one customer support, yet today, even with us, there can be hallucinations. So, while some clients are okay with an agent talking directly to their customers, we offer others human assisted support, which means a human is in the loop. We're also seeing significant interest in knowledge discovery and analysis across various sectors from financial analysts to legal professionals, and even in pharmaceutical research and media monitoring. Our current client base spans these areas, demonstrating the versatility of our product.

Some of our use cases include:

- Customer support: via digital agent chatbots or human agent assist, providing Q&A capabilities.
- Knowledge discovery, research and analysis: aiding financial analysts, supporting legal cases/contracts review, assisting pharmaceutical research, and enhancing media/news monitoring.
- Legal e-discovery
- Monitoring for insider trading/anti-money laundering (AML) compliance
- Analysis of user-generated content
- Enhanced eCommerce search functionalities
- Robust enterprise search solutions




Figure 4.1 – Vectara generative AI use cases

LISA: Will you specialize in certain verticals such as healthcare? And do you believe that domain expertise can be applied across different sectors?

AMR: Absolutely, our concept is truly elastic across industries, but to successfully launch a platform – or any startup for that matter – you need to focus on specific verticals. And to speak confidently within a vertical, you need to know the language. For instance, in healthcare, you need to know how to talk to medical people, and be conversant with medical terminology and compliance requirements. Similarly, if you're getting into finance, you need to have some compliance. So, we have four verticals that we're very focused on right now: health, legal, finance and media/tech. But there are also many others including manufacturing, customer support, and a smaller vertical for call centers. Further verticals will come later on, but these are the ones we're picking right now. We want to be an analytic platform that is useful for many, many verticals.

There will be large language models that are optimized for different verticals, for example, Google's Med-PaLM excels in the health sector, Bloomberg's finance-oriented model is outstanding, while Harvey targets legal applications

in London. Our capability is to federate across these models. We don't build the language models ourselves, instead, we position ourselves as the final integrative step.

Clients often have interdisciplinary queries. A hospital, while primarily concerned with medical questions, may occasionally need to address legal, accounting, or financial matters. Our platform is designed to intelligently detect the type of question and utilize the best language model for that inquiry, simplifying the process for our clients.

Our clients don't need to research which model to use, we handle it seamlessly on their behalf. This curation is one of the core values we offer, much like Twilio's facilitation of communication across various telecom networks or Stripe's streamlined access to financial networks, so you can charge a client money and have relationships with Wells Fargo and J.P. Morgan and Capital One. Or you can just sign up with, for example, Stripe once and they say, "We have a simple API and we're going to pick the best banking partner for you depending on what you're trying to do." Similarly, we arbitrate across the most effective resources, ensuring our clients benefit from the best available solution.

LISA: As a new product, what customer expectations do you have to overcome?

AMR: The biggest hurdle is the integration part, how to get the data from the client's content system into our system. This is where the friction tends to be. It's more about "how do I plug in the pipes in the right places?" But once you plug in the pipes, it's smooth sailing from there. That's my experience from many clients.

Key Takeaways

- **Context to generative AI:** The AI evolution is currently focused on generative AI, following a wave that included:
 - **Pattern recognition AI:** detected fraud, identified quality issues, and diagnosed diseases by recognizing data patterns.
 - **Clustering algorithms:** grouped product/user behaviors for effective marketing strategies.
 - **Forecasting:** predicted future scenarios, such as machinery failure or stock price movements, using time series analysis.
- **Impact of generative AI on software applications:** Generative AI deciphers human language and enhances the user-software interface. These technologies are expected to impact many software applications and improve efficiencies so that tasks that once consumed hours or days can be completed in mere minutes.
- **Learning process of AI:** The learning process of AI involves absorbing vast information and distilling it into a compressed form, much like human learning. The process of data compression and filtering involves an initial review and refinement by a human curator and post-training correction of misguided outputs.
- **In-depth industry knowledge:** A strategic focus on specific industry verticals, and the understanding of the unique language of each, is imperative for a more context-relevant LLM solution.



5

Philipp Heltewig

Co-Founder and CEO at Cognigy

There is no magic bullet. Implementing AI solutions in customer service still requires careful planning, development, time, and, crucially, human involvement...Looking ahead it's reasonable to predict that most customer service interactions will be handled by AI systems, except perhaps for very premium offerings. The adoption of AI by BPOs is a positive step toward a more efficient and cost-effective model, blending the consistency of AI solutions with human touch where it's most valuable, offering the best of both worlds.

LinkedIn: [heltewig](#)

PHILIPP HELTEWIG

Philipp Heltewig is a German tech entrepreneur with a strong foundation in computer science and business. He has international experience, primarily in the field of enterprise software. He previously led the global sales for Sitecore, an AI-driven digital marketing company, until its sale in 2016. Following that, he co-founded Cognigy, which has become a prominent figure in the conversational AI industry, focusing on automating contact centers for clients around the world. Cognigy's self-service solutions aim to improve the customer experience and support contact center agents in real time as they interact with customers. Philipp is recognized as a pioneering leader in creating conversational AI experiences and has authored content on innovating customer service through voice AI. His contributions at Cognigy have been pivotal in transforming both customer and employee experiences in the corporate sector through conversational AI.

LISA: How do your clients benefit from using AI in customer service, and what unique advantages do AI services provide?

PHILIPP: The primary objective for most of Cognigy's clients is to enhance customer satisfaction. While cost savings are a real potential benefit, they can be realized in various ways through implementing AI solutions in customer service.

The common perception of AI in customer service is largely centered around chatbots. However, AI-driven services can offer much more, such as real-time translation that enables customer services in multiple languages, even if the agents are only proficient in English.

LISA: How does Cognigy work with clients to maximize AI's potential in their operations?

PHILIPP: I'll describe how we collaborate with clients to identify needs. First and foremost, it's important to clarify that Cognigy is a software vendor. We partner with companies like Deloitte, Accenture, and Deutsche Telecom, who handle the bulk of the implementation of our product. Our clients generally fall into two categories. The first type is more experienced and understands

that conversational and generative AI is the solution to their problem, and they typically have their ROI (Return on Investment) models built out. The second type of customer is aware of the problem but may not fully understand the potential solution, such as implementing AI. It's this second group whom we often work directly with, or sometimes through a partner, to help them understand and establish ROI cases.

Our deployment strategy revolves around four key factors:

Customer satisfaction: This is straightforward. A satisfied customer is more likely to be a repeat customer.

Cost savings: These can be achieved through operational efficiencies in the contact center, such as improved translation or overall agent efficiency.

Improved agent experience: This is crucial in an industry prone to high agent turnover due to stress and burnout. Improving the customer experience also improves the agent experience, reducing negativity and turnover.

Increased revenue: This can be achieved through upselling. For example, we recently worked with a retailer of premium wines. By implementing agent assist technology, which listens to the conversation between the buyer and the agent and makes recommendations based on purchase history and customer profile, we were able to increase the average basket size, leading to increased revenue.

We work closely with our customers, often in workshops, to help them understand these benefits and build a business case to justify the investment in our technology.

LISA: How does AI integration shape your project stages and address challenges?

PHILIPP: Our key initiatives are the critical steps or objectives in a project. While AI is transformative, it doesn't alter the fundamental structure of our project stages, which include initial presentations, discovery workshops, and implementation planning. AI acts as an aid at various stages, enhancing ideation, assisting in project planning and content creation. However, due to particular

uncertainties, specifically related to compliance with the European Union AI Act, we allocate extra time to the resolution of these concerns.

LISA: Do you encounter misconceptions about AI's role in development and creation?

PHILIPP: I think it's essential to understand that AI, while increasing efficiency and simulating human tasks, is not a standalone solution. It's a tool that helps us get things done more efficiently. The belief that AI can independently produce a final product based on a simple command is a misconception. It may work in a basic demo scenario, but, in reality, human involvement is still crucial. AI is here to assist us, not replace us.

LISA: When and how does Cognigy approach vendor partners?

PHILIPP: We actively partner with leaders in the contact center space, including Genesys, Avaya, 8x8, and NICE. Some of these vendors have their own AI capabilities, which may be sufficient for some customers. However, for more complex projects with clients like Allianz Insurance or Lufthansa Group, their out-of-the-box capabilities may not be enough. This is not straightforward given the challenges of fragmented systems and legacy architecture, plus the need for data integration and regulatory compliance.

It's crucial to ensure that data access is compliant with regulations. As a company founded in Germany, we've had to deal with GDPR from the very beginning and we hold all kinds of certifications and have strong tools around data protection, data privacy, and compliance auditing.

In many large tenders, the list of potential vendors often gets whittled down to just a few due to compliance requirements. Our ability to meet these requirements is one of our core strengths and a key reason why we are a leader in enterprise conversational AI.

LISA: Given your deep capabilities, is your architecture and approach more vertical than the horizontal approach of larger vendors?

PHILIPP: This is a pertinent question. At Cognigy, we essentially offer two products. The first is our platform offering, which is designed for customers who want to build and control the AI models themselves. This is ideal for those who have established centers of excellence or are working with partners.

The second is a comprehensive solution we call AI as a Service (AlaaS) where we work closely with the customer to relieve their burden by understanding their AI needs, and then manage everything for them. This is a cloud-based offering that allows organizations to leverage AI tools and services without needing to build and maintain AI infrastructure and expertise in-house.

While our platform is not specifically verticalized, we do offer starter kits for various industries, recognized by Gartner Research as one of our core strengths. However, our AI as a Service approach is more verticalized, focusing mainly on five specific verticals where we have built up distinct intellectual property. This allows us to provide tailored solutions to meet the unique needs of different industries.

LISA: How are you mitigating the risk of deepfakes and generative AI being used by bad actors in customer service?

PHILIPP: The nature of our industry makes it difficult for deepfakes to be used maliciously so their impact is minimal on our business. However, the societal implications are indeed concerning. The ability of AI to create convincing deepfakes, whether in text, audio, or video form, poses a significant threat to the trust a society has in the authenticity of information to make informed decisions. It's a complex problem that requires a multifaceted approach to solve, legislation alone won't be enough to address this issue.

LISA: How do you assess the readiness of a company to implement your service and what prerequisites must be in place for a successful customer service implementation?

PHILIPP: The level of AI sophistication desired can lead to different scenarios. Here are three we often see:

Knowledge-based AI: This type of AI absorbs your PDF or website content and can answer questions about it. It's essentially a better type of search. For example, if you ask, "My dog Lily weighs 9 kilos. Can I bring her on board my flight with me?" the AI can retrieve the correct answer and generate a fitting response, such as, "I'm very sorry. Unfortunately, you cannot bring Lily on board because the maximum permitted weight for dogs is 8 kilos."

Concierge services: Sometimes called Identification and Verification (ID&V), this involves interacting with a bot that figures out what you want and who you are. It then hands you over to a human agent who is passed this information, reducing average handling time. Some back-end system integration is needed for this, to be able to authenticate the call.

End-to-end process automation: This is where you want to handle a full rebooking, refund, claims process, or package tracking, and so on. This requires deep integration into the enterprise's systems.

The success of these scenarios depends less on AI itself and more on the preparedness of the existing systems, including the robustness of enterprise infrastructures and the accessibility and efficiency of the processes they support.

LISA: How is the traditional model of outsourcing customer service during peak times evolving with the integration of AI technologies?

PHILIPP: It's a very good question because outsourcing to BPOs, especially overseas, has been common to manage service spikes, but AI is changing the game. AI offers benefits like multi-language support and instant scaling, which are reshaping how we handle customer service. While BPOs are adjusting to this trend, we're moving toward a hybrid workforce of AI and humans.

Looking ahead, it's reasonable to predict that most customer service interactions will be handled by AI systems, except perhaps for very premium offerings. The adoption of AI by BPOs is a positive step toward a more efficient and cost-effective model, blending the consistency of AI solutions with human touch where it's most valuable, offering the best of both worlds. It's an exciting time for the industry.

LISA: In what ways do customers manage setup and maintenance costs on their models, and how dependent are they on your support?

PHILIPP: The belief that AI solutions are always costly to set up and maintain doesn't hold true for conversational AI. Most companies often start with a core AI model and integrate their specific information, which is less about fine-tuning and more about building on a robust foundation. While they may initially require regular review and extension, they quickly (for example, within six months), become highly efficient, reducing the need for frequent fine-tuning.

Generative AI models have a vast amount of basic human knowledge built in, enabling them to provide sensible responses to a wide range of queries with minimal oversight. This highlights the power and potential of AI to streamline customer service, so that upkeep becomes less intensive over time, with our clients becoming less reliant on us for maintenance.

LISA: How do you manage the dynamic nature of customer knowledge, especially in industries like insurance where new classes of insurance emerge. And how do you address the cost and resources associated with retraining your AI on updated company collateral?

PHILIPP: The implementation of AI depends on its complexity level. For basic informational uses, updating AI is easy and efficient, for example, if a new product is introduced, the relevant PDF can simply be uploaded, and the AI quickly learns the information. However, this type of AI isn't designed for transactions, like processing purchases or modifying existing products.

Fully transactional AI is more challenging to implement, especially in insurance where there are few fully transactional cases. Checking a claim status is relatively simple, but processing a full claim is more complex. Some companies prefer to invest in smart routing and identification on various channels rather than tackling transactional use cases due to better ROI.

In the airline industry, there are high-value use cases like refunds or re-bookings that justify the time and effort to build out. For instance, Lufthansa Group has automated tens of thousands of rebookings and refunds during peak periods.

LISA: What is the timeframe from initial contact to the delivery of value within a company for this type of endeavor?

PHILIPP: Implementing an AI solution can range from days to months, depending largely on the type of functionality you want to roll out. For instance, if you simply want a bot that can answer questions about the content on your website, that can be set up quickly. However, the process can be lengthened by factors such as legal approvals, compliance checks, and technical implementation.

The challenge in rolling out a contact center automation solution isn't primarily impacted by the work with the AI, but rather external factors like the availability and accessibility of the right APIs, their documentation, and how ready the knowledge base is. In an ideal scenario, where all these factors are perfect, implementation could be completed in a week, but this is rarely the case. In large companies, issues such as acquiring API keys can take months.

The development of an AI solution, whether it's an FAQ bot or a bot capable of handling full transactions, is an ongoing process. For example, at one of our clients, their team went live with their first solution within two months but have since been working for four years on enhancing its capabilities so now the solution can do much more than it could initially.

The initial rollout can be quick, but the system needs to be extended and improved over time. This is akin to building a website for a new company. The first one-page website can go live within a week, but a comprehensive site

with various transaction capabilities can take up to two years to develop. The truth is the development process never really stops as there's always room for improvement and expansion.

LISA: Given this iterative process, it seems a nuanced understanding is crucial to set realistic expectations and achieve sustainable success.

PHILIPP: Absolutely. The advent of generative AI is a double-edged sword. On one hand, it opens up a plethora of amazing use cases, especially in our field, allowing us to deliver on the promise of human-like speaking.

On the other hand, the capabilities of generative AI are often overestimated by those not deeply involved in the field. For instance, while ChatGPT has indeed read the entire internet, it doesn't mean it can instantly provide all the solutions. It's a useful tool, but it won't dramatically change the timelines of an enterprise-level IT project.

In essence, generative AI is a valuable addition to our toolkit, but it's not a magic bullet. The implementation of AI solutions in an enterprise context still requires careful planning, development, time and, crucially, human involvement.

LISA: In your view, what percentage of the service gap has been closed?

PHILIPP: The impact of AI implementation can be quantified in various ways. For instance, E.ON, a large utility provider from Germany, achieved 70% automation on their conversational channels through the application of chatbots, and Mobily, a telco from Saudi Arabia, has a 99.5% faster response time due to their bot. Bosch achieved a 76% automated resolution rate in one of their conversational AI solutions.

However, quantifying the impact can be challenging for many large enterprises, especially those with unionized employees, as cost savings are seen to potentially lead to layoffs, even though we haven't witnessed that before. Therefore, these figures are often not shared externally.

The true measure of impact might be the potential disruption if the AI solution were turned off. If our customers were to turn off their conversational AI solutions, their contact centers would struggle significantly. This indicates that conversational AI has evolved from a “nice-to-have” to a “must-have” in many industries.

LISA: Could you elaborate on the specific generative AI models that Cognigy is using?

PHILIPP: Cognigy is an all-inclusive enterprise application that integrates all the necessary components for conversational AI applications. It incorporates LLMs, including OpenAI, Azure OpenAI, Google Gemini, Anthropic Claude, Cohere, and Aleph Alpha. The platform allows for the combination of these LLMs in a single conversation, providing maximum flexibility to its users.

LISA: Could you talk us through a specific customer project that you’re proud of?

PHILIPP: Certainly. A good example of how Cognigy.AI can deliver solutions in an agile way is in our work with Lufthansa Group. During the pandemic, we partnered with them to leverage conversational AI to transform their customer service in a time of critical need.

The Business Challenge

The core of Lufthansa Group’s business model is to bring people from one destination to another. They used a highly tuned model, which enabled them to forecast booking patterns and manage overbookings effectively. Cognigy had already started an initial integration of conversational AI with Lufthansa Group but the use cases were all based on business needs before the pandemic which, obviously, changed the paradigm.

The unprecedented disruption to normal life was acutely felt in the aviation industry. Lufthansa Group's service centers experienced a surge in call volumes with new types of queries and an increase in the average call handling time both due to their growing complexity and because junior staff were increasingly having to deal with them. This all resulted in service centers often being unreachable, which meant a negative customer experience that led to bad press. Lufthansa Group came to us to help tackle these challenges.

Context-dependent issues

- **Increase in manual contact volumes** when the business is operating at peak and there is an event in the network, such as a thunderstorm over one of the major hubs, the entire business is impacted. In the pandemic, this was acute with a surge in the volume of people contacting us over a couple of weeks, with all the stress associated with canceled flights and refund requests.
- **Higher average handling time of calls** as cases grew in complexity, multiplied by high staff attrition. More junior people were working on more complex cases which is the worst combination in customer service.
- **Increase in chat channel usage** as customer communication preferences shifted to use of chat channels through Instagram, direct messaging, Twitter, Facebook Messenger, and web chat.
- **The type of customer questions evolved** from "how many bags can I bring?" to baggage allowances to specifics about health and safety protocols, such as the type of masks required for flights from Germany to Dubai.
- **Poor accessibility of customer centers** during the COVID-19 pandemic, led to a deterioration in customer service.

Approach

Cognigy's agile conversational AI development was restarted from previous projects with Lufthansa Group to address specific customer needs during different phases of the pandemic's impact on business continuity, focusing on priority use cases:

- Re-open the skies was initiated with one airline, one use case and one language, focusing on providing travel information and requirements via a simple click bot.
- Offering refund capabilities due to the impact of virus variants on flight schedules. This transactional use case was widely used.
- Flight schedule changes enabled self-service rebooking for customers due to staff shortages disrupting flight operations.
- Managing peak operations by proactively handling customer re-bookings.

The setup included five teams developing the conversational AI experience for Lufthansa Group's customer service, with other business units joining in. These teams were:

- Continuous Improvement that focused on current use cases and FAQs.
- Rebooking that built two major rebooking functionalities – schedule change and voluntary rebooking.
- Personalization and Payment that managed the integration of customer profiles and personalization.
- Service Centre Integration that focused on integration with the service center.
- Brussels Airline Migration was a fourth airline that joined the Cognigy infrastructure.

What's next?

- Fine-tuning: We're gearing up for peak season by optimizing our operations to ensure a seamless experience for our users.
- User profile integration: We're enhancing the chat experience by integrating user profiles. This will eliminate the need for re-entering booking codes and names. Customers can just enter their details once in the app, and you're good to go!
- Service center integration: We're deeply integrating our services into the service center landscape, with a special focus on voice interactions.

Results and Impact

These solutions were rolled out across multiple airlines with significant and growing positive impact on levels of customer service.

Impacts included:

- There has been a significant shift from manual contact volumes to self-service platforms. Simple cases are now automated, resulting in shorter chat lengths, with the complexity of the case determining the extent of conversation required.
- Since 2019, chat channel volume has increased by over 2,000% and is projected to continue rising.
- The understanding of topics within chats has greatly improved, offering clear insights into dynamic changes and areas of competence. This enhanced clarity and transparency benefits both chat and voice platforms, with chat offering a more transparent view of what is effective.
- The system is highly accessible and scalable, managing hundreds of concurrent chats efficiently, contrasting with the limitations of handling a similar volume of voice calls.

- Machine translation has been employed to reduce language barriers, allowing agents to assist customers in different languages seamlessly. For instance, English-speaking agents based on Fiji can effectively handle chats in German for a traditional airline, with neither side being aware of the language differences, thus fostering inclusivity and improving customer assistance.

LISA: Could you please provide some insights into your discussions as an advisor to government ministers on AI-related matters?

PHILIPP: These discussions have been really insightful with some common themes emerging which I would emphasize to business leaders. The fact is AI is already here. It's not a future technology. It will touch all aspects of our business and professional lives so, if you're not focused on and thinking about its application, you need to be or else you will get left behind. I think every company should have a Chief AI Officer and every government will have a Minister for AI. Educating our children in the use of these technologies now is paramount. Other geographies, such as Asia, are already dramatically ahead. Finally, I think governments have an important role in tackling the issues of truthfulness, fake news, and so on to protect democracy.

Key Takeaways

- **Transformative impact of AI in customer service:** AI-driven services go beyond chatbots in customer service. They offer real-time translation for multilingual support, automate responses to customer queries, and provide real-time customer interaction data to agents. This enhances efficiency, reduces handling time, and creates upselling opportunities. The main goal is to boost customer satisfaction, with cost savings as a potential benefit.
- **AI in project stages:** AI doesn't alter the fundamental structure of project stages, but acts as an aid at various stages, enhancing ideation, assisting in project planning and content creation. However, AI is not a standalone solution, it's a tool that helps us get things done more efficiently.

- **Data access and compliance:** It's crucial to ensure that data access is compliant with regulations. This includes adherence to standards such as the General Data Protection Regulation (GDPR), obtaining necessary compliance certifications, and employing robust tools for data protection and privacy. Regular compliance auditing is also a key component of this process. These measures collectively help in maintaining the integrity of the data and safeguarding it against potential breaches.
- **Assessing company readiness:** The readiness of a company to implement AI services depends on the level of AI sophistication desired. This can range from knowledge-based AI to concierge services to end-to-end process automation. The maturity required for these scenarios is more about the readiness of the systems in place and the accessibility and processes around these enterprise systems.
- **Evolution of outsourcing model:** The traditional model of outsourcing customer service during peak times is evolving with the integration of AI technologies. AI offers benefits like multi-language support and instant scaling, reshaping how customer service is handled. The future likely involves a hybrid workforce of AI and humans.
- **Timeframe for AI implementation:** Implementing an AI solution can range from a day to three years, depending on the type of functionality you want to roll out. For instance, when implementing an automation solution for a contact center the factors that will most impact the timeline are related to preparation such as the availability and accessibility of the right APIs, their documentation, and how ready the knowledge base is.
- **Ongoing development process:** The development of an AI solution, whether it's an FAQ bot or a bot capable of handling full transactions, is an ongoing process. The initial rollout can be quick, but the system needs to be extended and improved over time.

- **The future of AI in the BPO industry:** AI is expected to increasingly handle most customer service interactions, signifying a major BPO industry shift. Knowledge AI can be easily updated, but in industries like insurance with few fully transactional cases, scaling the use of AI is more complex.
- **Mitigating risks in AI implementation:** Ensuring data access compliance with regulations, holding various certifications, and employing robust tools for data protection, privacy, and compliance auditing are crucial.



6

Miao Song

Chief Information Officer at GLP

With AI technology evolving apace, businesses need to ask how they should set up to support this technology, considering their layout, energy consumption, and future digital infrastructure. A robust new business model needs to think through all these things

LinkedIn: [miao-song-815877](#)

MIAO SONG

Miao Song is the Chief Information Officer at GLP and directs all facets of technology within the organization, and spearheads the company's worldwide data and digital strategy.

Prior to her role at GLP, Miao held a board member position and served as the global CIO at Mars Petcare, where she played a pivotal role in driving digital transformation and innovation. Before that, she was Vice President and CIO at Johnson & Johnson.

In 2019, she was named Chief Information Officer of the Year – Europe, by the International Data Group.

LISA: In your role at GLP, what was the core business problem you were trying to solve?

MIAO: Fundamentally it was to use data to create value – to help the company grow both their top and bottom line. And analytics was the core of this digital transformation – to drive more value from our data assets.

I focused on two components: how to integrate data science to create new value and how to increase business intelligence for customers and investors across the enterprise.

Until recently, there was a great deal of reporting on what had happened in the past and little about what is going to happen in the future. The business context contributed to this problem too as the whole landscape was behind the times.

Leveraging both internal and external datasets, we've now made good progress and launched a new business model across different regions. Throughout the process, we maintained a clear focus, asking for every initiative, "What is the business problem we are solving?"

MIAO SONG

LISA: Can you explain your initial process and what your first moves were?

MIAO: The process was challenging as it was a very decentralized organization, with lots of different initiatives. In addition, there was a lack of well-defined processes and structures, as well as local data governance and data ownership.

To identify potential AI applications, we began by examining several key areas of our investment practice: the assets, the properties, and the customers.

This analysis revealed the need for a comprehensive data management strategy. We implemented holistic governance and new technology to support the AI strategy.

One of the first things we also did two years ago was to hire new talent, which brought in fresh thinking and energy. Now, we still have a way to go but we have generated a lot of excitement about what we can deliver by using data better.

LISA: What leadership support did you have?

MIAO: We had strong global sponsorship from our CEO and leadership in each country which led to significantly improved visibility of our investments, allowing us to share them across the company more effectively.

LISA: What was the state of data when you arrived?

MIAO: The data was very fragmented, with quality and integration issues. So, we launched a data governance and Master Data Management (MDM) program which was a core foundation. We did a lot of hard work, such as diving into the property details, and the whole process was challenging and took 12 months to ensure it was correct and to get decent quality. In parallel, we started to collect internal data into a single platform and external data (for example, competitor data and marketing information), and inject it into the same data platform.

The company's previous variable approach across regions stemmed from the fact that their investment approach varied across regions, but with the rapid

globalization of the business we simply had to build a global data platform and build the property data asset into the single platform.

LISA: What changes and benefits are you seeing from the new systems?

MIAO: We've now got global visibility of the core business and enabled automation of the business insight work. This is enabling us to move from lagged reporting to prediction and forecasting. In the investment space, we're using data to make predictions – all things we just couldn't have done without the foundational work we did.

In terms of data processes, the nature of the investment industry tends to be more fragmented, compared to something like consumer healthcare, which does lead to quality problems. There's still a lot of work on modeling and capturing in Excel which we're trying to automate and digitize, but 50% is still being processed by analytics on Excel. Our go-forward approach has two key components: (1) digitizing data processing to streamline and improve efficiency, and (2) implementing data-driven modeling. We are using cloud tools but I'm realistic, so we're not getting rid of Excel but instead minimizing its usage and improving accuracy.

LISA: What resources did you put in place?

MIAO: We have a hybrid model between a decentralized and global team – and we also have a data team in China, Europe, and Singapore, too. We have created a new global data community of practice and co-drive initiatives.

We employed a consulting firm to bring in best practices, so we don't have to figure everything out on our own. And, for governance, we used a consulting firm to understand the methodology and mixture.

MIAO SONG

LISA: How did you maintain focus?

MIAO: We built through intention, only using business cases and KPIs with a focus on problems to solve. We challenged ourselves to continually look at the outcome of every single initiative and clearly forecast the size of the prize, for example, linking to the evolution of a global customer strategy.

LISA: Can you share any insights regarding the transformation of businesses into AI-led enterprises?

MIAO: In a previous company, I observed a large data initiative with a large investment really flounder because it lacked a clear objective. It took two years, and it yielded no business value. Just a change from A to B. This really illustrated the fact that you need to be business and not tech-driven if you want impactful results. Focus on tackling problems that will make a fundamental difference. So, I would advise not starting too big or with initiatives with long timeframes. Aim for small initiatives that make a big impact to show tangible results and gain top-level support.

Don't overlook fundamentals such as data quality and governance. Establish the community of practice on day 1, rather than a year later, and work with the business hand in hand. Finally, it's important to find good talent in the space you're operating in.

LISA: What are you most excited about in the future?

MIAO: Undoubtedly, it's OpenAI. With the technology evolving faster than we anticipated, we're asking what does it mean for our business when it comes to piloting and testing, and what can we do better?

With AI technology evolving apace, businesses need to ask how they should set up to support this technology, considering their layout, energy consumption, and future digital infrastructure. A robust new business model needs to think through all these things.

This tech is truly revolutionary, and we can already see in our industry how it's going to impact legal jobs, and fintech analytics where research can be made considerably more efficient. I know young people who were thinking about jobs in investment banking who are now reconsidering their choices.

Machine learning and AI are radically changing our business landscape with lots of investment going into this area. Every large company is working on it right now. It will change the way businesses will run, and public services, such as healthcare in particular, will change tremendously.

Key Takeaways

- **Foundational role of analytics:** Data analytics serves as the fundamental stage for any successful AI initiative. By implementing robust analytics practices, organizations can extract valuable insights to drive strategic decisions and actions. This analytics foundation prepares the organization for strategic AI applications.
- **Aligning AI with business goals:** To ensure impactful results, it's vital to understand the business context and implement a comprehensive data management strategy with governance. Rather than being solely tech-driven, this strategy should focus on business priority initiatives.
- **Talent acquisition:** Hiring new talent is a key step that can infuse fresh thinking and energy into the organization.
- **Leadership support:** Securing strong sponsorship from the CEO is crucial in inspiring the organization to be ambitious, prioritizing focus and investment, including redesigning the business model to make the best use of AI advancements.

- **Overcoming data silos:** A decentralized operating model creates fragmented data and integration issues. To address these challenges, a strong data management foundation is crucial. It should include data governance, Master Data Management (MDM), and a global data platform to integrate data assets from various sources.
- **Community of practice:** Establish a community of practice from the outset to accelerate cross-functional collaboration, problem-solving, knowledge sharing, innovation, and learnings. This builds trust and understanding.



7

Ruben Ortega

General Partner at Enjoy the Work

I was talking to a VC recently who just raised a \$20 million fund all focused on generative AI and I said, “What are you looking at?” And one of the first things they led with was “We’re looking for companies that have access to unique data.” It’s becoming a core part of their investment theses.

LinkedIn: [rubeneortega](#)

RUBEN ORTEGA

Ruben Ortega is a technologist and leader who has previously worked with Amazon, Google, Nordstrom, and Pinterest. He has built technologies and products that have scaled to millions of transactions per day, and generated billions of dollars in revenue per year. He established Amazon's first engineering office in Palo Alto, Google's Seattle office, and Pinterest's first engineering office in Seattle. He is currently mentoring startup founders as a General Partner at Enjoy The Work.

LISA: Give us a brief summary of your career up to now.

RUBEN: I graduated from the University of Washington in 1994 with an M.S. in Computer Science and Engineering and was fortunate to spend almost a decade at Amazon starting in 1998. I accumulated 40 patents related to data, coding, and technology, and built wildly successful search engine technologies. After that, I worked as the site lead for Google, focusing on search and their big query system. Following a brief stint as VP of eCommerce at Nordstrom, I spent two years building Pinterest's first distributed office in Seattle where we developed software in support of ads, marketing and machine learning. Today, I am an advisor to a variety of companies. It's exciting to work with startups and guide them in their growth.

LISA: What was it like in the early days of Amazon?

RUBEN: It was a really interesting time. The growth rate was incredible and it was truly amazing what we were able to deliver with so few resources in so short a time. I remember we delivered our first search engine capable of handling languages like German and French. However, we made some really amateur mistakes, such as neglecting diacritics, like umlauts, at one point! Luckily it mostly worked and as we got better, we had the resources to add the polish it needed to be great.

RUBEN ORTEGA

LISA: What are some tips for getting organizations to adopt data and AI practices?

RUBEN: I've discovered a useful approach is to reframe it as an experiment problem rather than a data or AI problem. Google did loads of great research that showed that psychological safety is crucial for effective teams. So, by changing the language from success/failure to control/experiment, we created a safe environment for discussions. Communicating that "failures" are part of the process and emphasizing ongoing experimentation can also help shift the mindset.

My question to any company is always "What experiments are you running and how are you gathering your data, then how are you storing it?" This is especially important when working on machine learning and AI. "Garbage in, garbage out," still applies. It's important to capture the right data, store and curate it properly, and iterate through pilot experiments.

LISA: How do you get the support to do this work and get it prioritized?

RUBEN: It has to show signs of success from the beginning. You can't boil the ocean and expect a company to pay for it to get that little one percentage lift along the way. Aligning data projects with the company's North Star metric, such as customer happiness, rather than revenue, is essential.

One example of this process was my first big patent at Amazon 18 years ago. It was a search relevancy feedback loop. I implemented it over a weekend, tried it on my desktop, then showed it to my boss and we both went, "Wow, that's... amazing!" We didn't know what it was but it was materially better than any other results at the time. So, with his support, we spent another weekend on it and then showed it to our VP who looked at it and said, "Can we do this again?" We filed a patent, spent another couple of weeks working on it, and then launched, and it was immediately credited with an important lift in daily revenue results. And that was when my VP had to go to [Jeff] Bezos and say, "All those plans I just launched for the year, we're gonna have to scramble them because we just discovered something." So that was a case of investing just enough to decide if we should continue investing. But it's illustrative of the

success of sponsorship stories. You have to show something a little bit exciting using the data to get people on board.

Examples of data that we know worked were every time we were able to use the purchase signal at Amazon. And Pinterest also kept looking for ways to capture more signal. The joy of Pinterest is that it was a visual discovery engine, and it grew up after Amazon and Google, so it already had experimentation built into the bones of the product.

LISA: How do you engage people who are resistant to the benefits of data and AI?

RUBEN: Sharing success stories and demonstrating how data and AI can significantly impact the company's goals are good starting points to generate energy and trust.

I've found a couple of effective strategies. Firstly, teaching teams to tell compelling stories and linking that narrative to the customer as the hero (note: the hero is not the engineer or management!), can make a big difference. For instance, I have an idea the data is telling me we have a certain ICP (Ideal Customer Profile) in a pharma company. So, tell the story, using the vocabulary of the business, "Here's the world as it is today, here's the crisis our customer is going through, here's how we're going to fix it, and here's the outcome for our hero – the customer. And here's a percentage lift we're going to see." Personalizing the story is crucial because you're dealing with humans who want to look good in front of their bosses.

Secondly, aligning data projects with the company's North Star metric, such as customer happiness, is essential. It's worth noting that revenue is not a good North Star because it's a lagging indicator so, for most data and AI projects you need to find the leading indicators and see which ones best get you to that outcome. For instance, I can tell you that \$5 billion is better than \$4 billion in revenue, but I also know the way to get there is by having ten thousand customers logging in versus five.

LISA: What do you do when you go into a company and their data is a mess?

RUBEN: The first step is cleaning up the data by throwing out the garbage. The reason why generative AI has become all the rage right now is that they've actually been pulling from high-quality data sources, but now they're having to filter out sources that were too high quality and were causing problems with copyright.

Depending on the type and importance of the data, different strategies can be applied. Setting clear definitions, such as that of the customer, and raising the bar can improve data quality. For instance, Pinterest's definition of a customer used to be somebody who'd come to the site in the last 90 days. However, this definition proved to be ineffective. Consequently, we refined our definition to include only those individuals who not only visited the site at least three times within the past 90 days but also engaged in certain activities during their visits. This temporarily dropped the number of customers they reported but, with a higher bar for customers, we could demonstrate a more accurate leading indicator to revenue.

Identifying the key metrics aligned with the company's North Star and ensuring accurate data collection, storage, and processing are crucial. For instance, say you have a company that doesn't know when to sell their product. They need to look at the leading indicators from previous sales. When you find those seeds of data, you do a post-analysis and start making sure all the incoming data is correct and say, "Okay, here's the golden three or five metrics." Once you learn the discipline of getting those in place, you can start adding the 6th, 7th, 8th... to test queries and make sure the data propagates through the system.

When you've identified the key metrics you want to track in place, you need to add the muscle within the organization, by adding more metrics – maybe one a month. Gradually building the necessary infrastructure and processes in this way ensures a sustainable approach.

RUBEN ORTEGA

LISA: How do you foster effective teamwork?

RUBEN: Ownership is vital. Designating a leader who takes ownership of a problem for a specific time period, with the resources needed for a solution, can be effective. Timeboxing the work also helps maintain focus.

Another crucial factor is aligning with the company's values, getting this right can lead to it becoming a scaling tool. Having the CEO say, "Hey, I just want to acknowledge this person taking up this problem. He really exhibited the value of X." Highlighting individuals in this way fosters a culture of teamwork and learning throughout the organization.

LISA: What lessons have you learned from situations that felt like failures at the time?

RUBEN: In the early days of Amazon, we had projects like Amazon Cars – a car referral search – that ultimately failed because the customers didn't click through and, when they did, they were confused. However, as we went through each iteration, the technology got better and faster, so each failure led to learnings that contributed to future successes.

The advice you give CEOs and founders is – you'll never have a perfect and happy team. I've worked for 30 years in software development and there's always somebody who is just learning the technology, bored of their work, getting a divorce, having a health problem, or a project manager who wants to be a software engineer and vice versa. But with clarity and effective management, you can always deliver fantastic things incrementally.

LISA: What's your experience with change management?

RUBEN: It depends on the company's values and desired approach. Iterative or a "pull off the Band-Aid" method can both work, provided you're consistent with what the end state is.

RUBEN ORTEGA

A Band-Aid example at Google was when we had to switch 300 people from working on search to working on something else. That's when you get in front of the crowd and openly communicate the plans, reasons and expectations. You tell everybody what you're doing, where you're going, what's happening, why it's happening, and then you execute.

An example of an iterative approach at Pinterest was when we were developing a studio model for developing features. We followed a pilot and evaluation approach, gradually rolling out changes and aligning with their collaborative values. So, instead of having the whole front-end and back-end team do it, we piloted it with a small team of seven who were accountable for this one feature and execution, including whether they close it down, spin it up, or grow it.

Defining a clear plan, seeking feedback, and iterating, helps navigate change effectively.

LISA: What are your reflections on the internal use of data and AI within companies? What works well and what unlocks it?

RUBEN: Marketing technology is a complex space because everybody has a different idea of what a strong signal looks like. To effectively use data within companies, marketing and sales teams should adopt a disciplined and experimental culture.

Implementing frameworks and tools like Hubspot and Copper are fantastic to assist in data-driven decision making but they need focus from the leadership of those organizations to be fully effective.

I'm an adviser to one company that does churn analysis and their tooling is great and they use AI so that you can take a look at how many contacts you had per month, etc. The big signal about churn analysis is not when the customer's contacting you, but when they stop, because they're saying "This product is not serving me." So, if they're not using the product and they're not calling you, you're dead. That's a churn. That's one signal and this company does that pretty damn well because they use AI and all the magic of it.

RUBEN ORTEGA

LISA: What are you most excited about now in terms of using data and AI within companies?

RUBEN: The increased discipline and knowledge around data utilization and AI are exciting. There used to be the phrase “Software is eating the world,” but now it’s data that’s eating the world and companies know they need to capture it, make it crisper, and leverage it for better decision-making. The hype around generative AI is fantastic but it again highlights that data is the thing.

I was the CTO of the Institute of Artificial Intelligence for a short time to help them grow their teams, but one of their theses of development was “How can we use the least amount of data to make an intuitive leap?” Making intuitive leaps with AI using smaller, high-quality datasets is an area of interest. It’s not just about flooding models with vast amounts of data but extracting meaningful signals from smaller datasets.

Investors are also prioritizing companies with access to unique and valuable data. I was talking to a VC recently who raised a \$20 million fund focused on generative AI and I said, “What are you looking at?” And one of the first things they led with was, “We’re looking for companies that have access to unique data.” It’s becoming a core part of their investment theses.

Key Takeaways

- **The power of unique data:** In the venture capital world, access to unique data is increasingly becoming a core part of investment strategies.
- **Momentum with early wins:** Initial results of AI projects should demonstrate early success and align with the company’s key metric. Showcasing AI’s impact on company goals can build trust and enthusiasm. Identify leading indicators for desired outcomes in data projects.

RUBEN ORTEGA

- **Storytelling and personalization:** Teaching teams to tell compelling stories and linking that narrative to the customer as the hero can make a big difference.
- **Experimentation:** To gain buy-in, frame your AI project as an experiment, not a data issue. Companies like Pinterest have shown that integrating disciplined experimentation into product development leads to innovative outcomes. Even failures provide valuable insights for future success.
- **Handling messy data:** To enhance data quality, establish clear definitions for core terms used within the business, for instance defining what is meant by customer. This ensures consistency and accuracy in data interpretation and analysis. Identify key metrics in line with the company's primary goal. Ensure accurate data collection, storage, and processing.
- **Fostering effective teamwork:** The best-assembled teams benefit from a dedicated leader with ownership and resources for execution within a clear timeframe. Ensure the project aligns with company values and recognize those who embody them to promote a teamwork culture. Embrace continuous improvement with ongoing communication and adaptation.
- **Change management:** Either iteration or a “pull off the Band-Aid” method can work, provided you're consistent with what the desired end state is. Defining a clear plan, seeking feedback, and iterating, helps navigate change effectively.



8

Joshua Rubin

Principal AI Scientist at Fiddler AI

We learned that the actual need in the enterprise extends beyond the explainability of the AI decisions to the broader story of operational consistency and risk management. So, our focus shifted from a purely technical perspective to one that encompasses standardization, centralization, and compliance. After companies deploy their models into the real world, they require extensive measurement and performance tracking to ensure that things are working as intended.

LinkedIn: [joshua-rubin](#)

JOSHUA RUBIN

Joshua Rubin is the Principal AI Scientist at Fiddler AI with a PhD in Physics from the University of Illinois. As a scientist and machine learning practitioner, he applies his expertise to algorithm and product design, specializing in model explainability, semantic drift in unstructured data, and model-based evaluation for workflows involving large language models. Rubin champions clear, explainable AI, integrating this ethos into Fiddler's strategy and product suite. Leading Fiddler's data science team, Rubin balances technical skill with strategic foresight, aligning product capabilities with client requirements. His pragmatic approach ensures Fiddler's offerings not only meet current needs but also adapt to future demands. In an industry prioritizing clarity and utility, Rubin focuses on delivering accountable, user-oriented AI solutions.

LISA: Can you tell us about the background of Fiddler AI?

JOSHUA: Fiddler is an enterprise AI observability company founded by Krishna Gade and Amit Paka. They founded the company to address the critical need for AI observability and explainability in the field of machine learning and artificial intelligence. Their vision has led to the development of a robust platform that empowers organizations to monitor, analyze, and improve their AI models while ensuring transparency and accountability. Before founding Fiddler, Gade held an engineering manager role at Facebook, leading a team that built tools to enhance transparency in the product's News Feed personalization. His experience at Facebook, especially after the 2016 US election, made him realize that the challenges related to the machine learning workflow – such as monitoring, explaining, and maintaining models – were not unique to Facebook but were widespread. This realization led to the foundation of Fiddler in 2018 with a vision to create a unified platform for managing machine learning models, ensuring transparency, and addressing issues of fairness, accountability, and trust.

LISA: What are the key offerings of Fiddler?

JOSHUA: Fiddler provides an AI Observability platform, enabling companies to monitor, explain, analyze, and improve predictive and generative AI models. We offer a full stack platform that detects issues as soon as models underperform, measures and tracks ML and LLM metrics, provides rich insights into model outcomes using human-readable explainable AI (XAI), diagnoses model issues with root cause analysis, and assesses for fairness and bias. Through our responsible AI framework, we help companies build trust in AI.

LISA: Can you explain the importance and challenges of explainability in AI and how it impacts various stakeholders?

JOSHUA: In the realm of AI and machine learning, explainability is a critical component. It's not uncommon for end users to question why they're seeing certain content in their newsfeeds or why specific ads are targeted at them. These queries arise because of the highly personalized nature of these algorithms tailoring content based on individual behaviors and preferences.

The focus of explainability is to demystify the decisions made by these models. This need for understanding isn't limited to end users; it extends to model developers, technical overseers of AI behavior (MLOps roles), and executives who need to be able to trace and justify these algorithmic decisions, a process that becomes even more crucial for executives who must ensure these decisions align with regulatory standards.

Consider the banking industry: heavily regulated and often torn between interest in the financial benefits of using sophisticated AI for credit underwriting and the need to comply with transparency regulations. Traditional numerical models, such as credit scores, offer clear explanations for decisions, which can be in contrast to the often opaque nature of complex models, which poses a significant challenge.

Another example is the use of an image classification or object detection model in a model being used to evaluate insurance claims. Explainability can provide the “why” behind an AI recommendation, allowing a human domain expert to, for example, tell the difference between a fraudulent claim and a false positive produced by background imagery captured under conditions outside of the model’s training experience.

LISA: How do you apply explainability in situations of high speed and volume where AI algorithms are making decisions?

JOSHUA: Online advertising is a good example of this scenario. In this fast-paced domain, algorithms make split-second decisions on ad placements based on a myriad of user aspects. These rapid decisions, while impressive, are often opaque.

So, Fiddler has a dual role to play: informing recipients (like those seeing an ad or receiving a loan decision) about why they’re seeing what they’re seeing, and providing an offline diagnostic review or incorporating human-in-the-loop review, where specialists analyze algorithmic decisions, ensuring AI’s processes are auditable and can be refined to improve AI performance over time. These insights can influence a company’s strategic direction and compliance, as articulating AI’s decision-making is critical for maintaining regulatory compliance, particularly in sectors like advertising where data usage and consumer privacy are of utmost concern. Fiddler workflows can support both real-time explanation and retrospective analysis.

LISA: What type of customers do you sell to and what are their pain points?

JOSHUA: Our ideal customers range from hyperscale startups to enterprises that have built applications on top of ML models or LLMs. This includes banks using lending or fraud models, and cloud-native commerce companies whose algorithms determine the order of items displayed on a user’s screen.

Understanding the performance of these algorithms, especially when they underperform or when the world changes in a way the model wasn't prepared for, is crucial. This concept is often referred to as "model drift."

Take, for example, a credit underwriting model at a bank; imagine if interest rates were to change. This could not only alter the distribution of household incomes of loan applicants, causing the model to make unsafe extrapolations from its training data, but it could also change the real relationship between an applicant of a specific profile and their probability of successful repayment. Such a change can cause the model to perform less accurately than at the time of its training, potentially leading to lost revenue and risk. We've shown our customers examples of reputational harm and unnecessary financial losses, all resulting from models making poor decisions due to unanticipated changes in the world.

To mitigate these risks, we offer a pre-emptive solution that alerts model owners when their models may be underperforming or when even subtle change in the world might affect their performance. This allows for timely adjustments and improvements, ensuring the models continue to serve their intended purpose effectively and safely.

LISA: Can you explain in more detail how Fiddler tracks model drift?

JOSHUA: For structured data, Fiddler detects model drift in distributions of model inputs by periodically comparing histograms gathered for models in production with those gathered from training with a variety of standard statistical formulas. For models that consume unstructured data like natural language or images, Fiddler compares multidimensional distributions of embeddings – essentially numerical vector representations of this data. If drift is identified, the platform can notify model operators via email or other mechanisms and provide a variety of tools for root-cause analysis.

JOSHUA RUBIN

LISA: How would you say that Fiddler upholds the principle of responsible AI?

JOSHUA: Fiddler was founded on our founders' previous experiences by making model observability tools accessible to a broad range of stakeholders within the enterprise environment and throughout the AI lifecycle.

As conversations with customers evolved, it became apparent that Fiddler fits into and addresses challenges in the customer's process and approach. Once companies have built their models and have left the training zone to be deployed into the real world, they need a lot more measurement and behavioral performance tracking to make sure that things are working as intended.

You need to ask, how much measurement is needed? What will happen if they don't track or use AI observability? How will it impact them? What are the implications if they were to use an observability tool? How much more confidence would be instilled in them to push more models into production?

Fiddler has a responsible AI framework that is divided into four phases, where each phase represents growth in the AI journey, and milestones are mapped to business KPIs and ROI. AI-forward organizations view this framework as a way to rethink their orientation around change management as it relates to people, process, and technology. The principles of the framework also serve as guidance for the enterprise to be an AI-forward company for the long term.

The overarching principle at Fiddler is to build trust into AI. Whether it's monitoring, explainability, analytics, governance, fairness, or bias; everything that Fiddler can assist with is underpinned by the principle of responsible AI.

LISA: As you referred to very different sectors – advertising, social media and banking – how important is industry domain expertise for your services to benefit an organization?

JOSHUA: Our tooling is general so that it can be applied across various domains. This approach, to not target a particular vertical, was a business decision that Krishna and Amit made early on because they believed there was a substantial need and enough commonality across different domains to make it viable.

A key component of this strategy is to assist human domain experts, who often aren't the model developers. For instance, we worked with a talent placement company that had a model to evaluate resumes for suitability for roles where other companies had job listings. The goal was to aid career counselors in matching the right people with the right roles.

In this scenario, the career counselor, who is a domain expert, is assisted by an algorithm that provides a recommendation. If the algorithm can provide more information about the guidance it's providing, it can really help the human domain expert. For example, if the algorithm can explain why it thinks a candidate is a good fit for a role, or when a particular prediction may have fallen into a blind spot where the model doesn't have much training, it gives human domain experts extra insight into why the algorithm is making the recommendations it does.

This essentially supercharges the human domain expert. Humans have limitations, such as fatigue from looking at numerous cases in a day. Imagine, for example, being a physician working in a hospital setting, and trying to maximize the quality of care provided. Because a computer never gets tired, it can consistently provide important insights. This reduces the likelihood of human error and allows the human expert to work more efficiently.

Furthermore, when the human expert discovers that the model is behaving in a way that's not aligned with their expertise, it gives them a better ability to communicate this to the model developer, who may not have the same level of domain expertise. This collaboration between human experts and machine learning models can lead to more effective outcomes.

LISA: Based on the evolution of your product, how do you envision the future of AI in enterprise environments, particularly in terms of real-time monitoring and adaptability across diverse data types?

JOSHUA: It's a journey of discovery and adaptation. Initially, we were all about the technicalities, the nuts and bolts of AI models. But as we delved deeper and engaged with enterprises, it became evident that there's more to this narrative. It's

not just about how these models are crafted or why they make certain decisions; it's also about their integration into the broader enterprise ecosystem. We learned that the actual need in the enterprise extends beyond the explainability of the AI decisions to the broader story of operational consistency and risk management. So, our focus expanded to encompass workflow standardization, centralization, and compliance – elements crucial in regulated industries like banking.

Enterprises aren't just looking for tools; they're seeking solutions that streamline processes, enhance efficiency, and ensure compliance. Our product scope has changed in response to these needs – from developer tooling to an enterprise platform offering standardized observability layers. This evolution came from recognizing that enterprises often require a standard platform and process to avoid pitfalls of decentralized manual data entry and calculation, which increases risk from human error and can lead to regulatory problems or worse.

The challenge isn't just in creating AI models but validating them in pre-production and monitoring them in real-time production environments. It's about ensuring they continue performing optimally and adapting to changes in data inputs and environmental variables. The real need that emerged was for repeatable tests and having a software product that provides a standardized observability layer to larger enterprises.

This understanding was accompanied by a shift to address not just simple structured data, but also the more complex unstructured data types like natural language and image – especially with the advent of generative AI. Real-time monitoring becomes pivotal here; you can't wait for quarterly reviews to realize something is amiss.

So yes, our journey has been one of evolution – from focusing on algorithmic explanations to addressing real-world enterprise needs around standardization, compliance, real-time monitoring, and adaptability across diverse data types. It's crucial for businesses to have real-time insights into their operations. We have to think about who's watching models in production and how often they need to be checked to ensure they are performing as initially designed.

For instance, if a camera or sensor on an assembly line starts to fail, it's not ideal to wait until the end of the quarter for a review. By then, you might discover that you've been producing faulty products for the last two months. This highlights the need for a model monitoring product that tracks performance changes and data drift – how the world is changing over time and differing from the training data.

This is where automation comes in, with alerting and dashboarding features. A model operator can then monitor a portfolio of complex models, ensuring they're performing as designed at any given moment. It's not just about having models that perform some business function, it's about maintaining clarity on what it means to put a model into service and ensuring there's a system in place for continuous monitoring and adjustment.

This kind of value proposition is particularly appealing to our enterprise customers. They appreciate the ability to promptly address issues, maintain quality control, and optimize their operations. It's all about staying proactive rather than reactive. Our enterprise buyers are looking for the kind of value proposition that emphasizes immediate knowledge and action ensuring that any disruptions in the production process are identified and addressed promptly.

LISA: How does this new concept, as both a platform and a layer, fit into a company's existing tech stack?

JOSHUA: Fiddler provides a web interface where users can interact with charts, dashboards, and analytics, and programming interfaces for publishing data from your model inference pipeline. There are programmatic hooks that get stitched into the company's serving platform to publish to it. Users can configure alerts on different kinds of signals that might be indicative of a need for action. There are three main deployment options:

- **Software as a Service (SaaS):** In this mode, we host the Fiddler platform on the cloud and fully manage its maintenance and scaling.

JOSHUA RUBIN

- Virtual Private Cloud (VPC): Most companies with privacy concerns have a VPC with AWS, GCP, or Azure. These big hyperscale companies have a private network within the envelope of a cloud provider. All of their infrastructure resides virtually within that, and we can deploy a version of Fiddler within it. It basically becomes a sort of standalone scalable system to which they publish their data.
- Private data center: Enterprises that have very strict security policies can deploy Fiddler in an air-gapped setting.

The support demands of these three different deployment models are quite different. It's easy for us to support the SaaS platform as we can control scaling, see where there are problems, and upgrade it with our own hands. On the other end of the spectrum, with the private data center deployment, we never get to touch the keyboard and we're lucky if we get a screen share.

It's important to note that we understand there are different customer deployment requirements. Some customers send us a few events per hour, but then we have some that are ad tech or internet marketplace companies where it could be millions of events per hour. Building for scale was something that took a real effort. It wasn't the way the original platform came out, but we've crossed that hurdle and built something that really scales and is stable. It's pretty great, actually.

LISA: Looking to the future of the company, what are your thoughts and what excites you the most?

JOSHUA: The future of the company is incredibly promising, especially in the realm of LLMs and LLMOps. Last year, around October, we introduced metrics that allow you to evaluate how your models are performing, including their responsiveness, robustness, correctness, privacy, and safety. You can even assess if your models are hallucinating or producing toxic responses.

These metrics we've developed are proving invaluable to companies. As more initiatives and pressures trickle down from the boardroom into the core organization, it's crucial to stay on the cutting edge. If you're considering deploying chatbots, for instance, you need to ensure they're not toxic, abusive, or hallucinating. They must be safe. We've only scratched the surface, and there's so much more to come.

In addition to out-of-the-box metrics, companies can also measure specific metrics for their unique use case by creating their own custom metric by bringing in their own formula. Say you recently launched a chatbot and want to monitor the associated cost. You can create a custom response cost metric to track each OpenAI API call from the chatbot. And a separate custom metric to track the rejected responses generated by the chatbot – responses that end users have deemed poor or unhelpful. By combining these two metrics in Fiddler into a single formula, you can then calculate the true cost of using the chatbot. This approach allows you to gain a holistic understanding of the true financial impact and effectiveness of the chatbot.

AI observability is a massive field with layers upon layers and is necessary to ensure models and LLM and GenAI applications behave and perform in an intended and healthy fashion. At its core, it's built on a vast foundation of data. There are countless opportunities for growth, both downward into the data and upward into new applications. I can't even begin to imagine what we'll be able to do two to five years from now.

LISA: What advice would you give non-technical business leaders on how they might think about integrating AI into their companies? How does this relate back to your own company?

JOSHUA: I often tell my team and others that every technology has its positives and negatives. It's all about how you use it. I'm proud that our company has always prioritized building trust in AI and ensuring that our outcomes are safe and fair for consumers. This commitment to social responsibility sets us apart from many other tech companies.

JOSHUA RUBIN

Any new technology can be intimidating. There's a common fear that AI will replace jobs. AI will certainly change jobs. Those who don't adapt could find their roles becoming obsolete. This fear is not new; it was there when spreadsheets, telephones, books, and newspapers were introduced. It's all about how humans react to change.

Businesses that choose not to embrace AI will find that their competitors take advantage of that and outpace them. AI adoption is growing in organizations, with more applications using AI for a variety of purposes. Tooling, like we provide for AI Observability, is crucial, whether that's for larger language models or more traditional machine learning. And while there's some overlap, the two types of models are mostly complementary, solving different kinds of problems. Traditional machine learning continues to grow and remains critical for a wide variety of predictive tasks.

When you have a company with one or two AI use cases and a small development team, you can manage those things reasonably well. However, as you adopt more AI models and they function in increasingly mission-critical capacities, having tooling like Fiddler becomes essential. It provides model governance, protection against reputational harm and revenue loss, and allows you to adapt rapidly to changing conditions in the world.

There's also the centralization consideration. When you have a diverse model portfolio and a large team of developers, centralizing both MLOps and LLMOps in a platform like Fiddler provides operational efficiency. You don't want thousands of people in your organization manually rolling up statistics in Excel spreadsheets. You want a centralized algorithm that's been vetted.

LISA: Can you talk through a company that went through this standardization, particularly in observability, and scaling challenge? How did you work with them?

JOSHUA: Certainly. Let's discuss Tide, a leading UK business financial platform. The complexity of the company was increasing, so they came to us for this type of solution. Tide recognized the urgent need to scale their machine learning (ML) solutions to offer automated bookkeeping and integrated invoicing for small

businesses without a dedicated finance function. However, the maintenance of production ML models was a significant bottleneck, preventing their data science team from pursuing new business opportunities.

Business Challenge

Tide's data science team was inundated with the task of troubleshooting model performance issues. Due to a lack of standardization, each team member was using a different monitoring tool, leading to inconsistent views on model performance and creating decision-making silos across data science, machine learning, and business teams.

Context-dependent issues

- **Reactive model monitoring:** The team's model monitoring approach was so reactive that issues would go unnoticed until they became critical, leading to a longer mean time to detect concept or data drift, challenging root cause analysis, and complex model explainability.
- **Lack of standardization:** With each team using their own tools, there was no unified perspective on model health, which affected coherent decision-making and innovation.
- **Resource misallocation:** Data scientists spent excessive time managing scripts for troubleshooting rather than innovating and unlocking new business avenues.

Approach

To address these challenges, Tide's Head of Data Science, Suryanarayana Ambatipudi, recognized the need for an enterprise-grade model monitoring tool that could bring together all the teams on a unified platform, streamline processes, and enable scalability in line with the company's growth strategy.

After considering the development of an in-house tool, Tide opted to buy an external solution for cost-effectiveness and standard adherence. Fiddler was

chosen as the AI observability platform to meet their needs for a robust and flexible MLOps framework.

Results and Impact

The implementation of Fiddler as Tide's AI Observability platform led to substantial improvements in both operational efficiency and strategic focus.

Impacts included:

- **Higher productivity:** Tide's data science team experienced a significant increase in productivity. By utilizing Fiddler, the team no longer needed to dedicate extensive periods to monitoring and troubleshooting performance issues, which freed up time for value-adding activities.
- **Cost efficiency:** The automated performance tracking capabilities led to a reduction in both infrastructure and maintenance costs, marking a direct cost saving for Tide.
- **Strategic focus:** The team could now rapidly identify and resolve issues impacting the business. This agility enhanced their ability to provide swift post-mortem analyses, which is crucial for continuous improvement and risk mitigation.
- **Unified performance metrics:** A unified view of model performance metrics, ensured consistency in the evaluation and understanding of ML models across the organization.
- **Improved collaboration:** The improved collaboration between the data science and business teams facilitated more informed decision-making and reinforced the role of data science as a pivotal contributor to Tide's business strategy.

These results underscore Tide's commitment to leveraging cutting-edge technology to enhance their operational efficiency and maintain their position as the UK's leading business financial platform.

“Fiddler’s AI observability platform has enabled us to shift our conversation from merely being model providers to becoming confident decision-makers, aligned with business strategies and insights. This has transformed how we operate, and we anticipate a significant positive impact on our business metrics.”

— Suryanarayana Ambatipudi, Head of Data Science at Tide

Future Plans

Tide aims to integrate explainable AI within their framework to provide operational insights, thereby reducing the mean-time-to-resolve (MTTR) business metrics and aligning model outcomes with business decisions.

Key Takeaways

- **Mitigating risk through responsible AI:** Building trust in AI through responsible practices is no longer optional. It sets companies apart, safeguards against reputational risk, and protects against revenue loss.
- **The importance of standardization and observability:** Standardization and observability are crucial principles in responsible AI. These practices are essential for enterprises when building, validating, and deploying models into production.
- **Integrating AI into enterprise environments:** Enterprise AI goes beyond model creation and decision-making to integration into existing systems. This prioritizes explainability, operational consistency and risk management, standardization, and centralization.
- **Transitioning to AI platforms:** Enterprises are transitioning from developer tooling to an enterprise platform that standardizes observability layers. This allows for simplified compliance and real-time monitoring of the AI model performance for issue detection.

- **Pre-empting model drift:** When algorithms underperform or the world changes in a way the model wasn't prepared for, this "model drift" can lead to lost revenue and pose risks. These can be caused by a number of factors, including economic shifts, new customer behaviors, and seasonal variations. This can lead to less accurate predictions because the model is trying to apply outdated rules to a new situation. To mitigate these risks, consider a preemptive solution that alerts owners, so they can make proactive adjustments and ensure ongoing optimal model effectiveness.



9

Nadine Thomson

**Global Chief Technology Officer at
GroupM (WPP)**

Recently, I was looking at a product that can take a picture and modify the background based on user commands, while avoiding copyright issues. This is something that would usually take a long time and, although it still requires some tweaking from a brilliant creative to pull it all together, the speed of change is both daunting and exhilarating.

LinkedIn: [thomsonnadine](#)

NADINE THOMSON

Nadine Thomson is the Global Chief Technology Officer at GroupM, the world's leading media investment company. She has helmed B2B and B2C international transformation programs across the media, retail, insurance, professional services, and entertainment sectors, including leading technology for Conde Nast International and Vue Cinema, as well as advising News UK and Capita on technology strategy.

LISA: Could you tell us a bit about your journey in the technology and data space?

NADINE: I currently serve as the Group CTO at GroupM, the media investment division of WPP. My primary responsibility is to unify the technology across the holding companies within GroupM, put our money behind the best bets, and drive innovation. As you become more senior in technology, the job is often less about dealing with technology challenges and more about addressing people issues.

My journey in technology started in quite an unorthodox way as I initially went to university to be a vet but then swapped to a computer science degree. Since then, I've worked in a range of organizations including a stint in publishing with Conde Nast, then News UK, cinema, and oil and gas insurance.

Aside from my current role at GroupM, I'm also a non-executive director at Visit England, a position that helps me broaden my perspective on business and its key considerations.

Although seemingly diverse, in the various roles I've held, data and AI have become an increasingly significant focus. Whether working alongside or sometimes leading chief data officers, I find myself continually integrating and strategizing around AI as it is integral to all of our products.

NADINE THOMSON

LISA: Could you delve a bit deeper into your approach to data and AI when joining a new organization? What has been the problem that you're trying to solve and how did you get the sponsorship to move that forward?

NADINE: Whenever I join a new organization, I don't necessarily come in thinking "What is the tech or data challenge?" but instead I focus on understanding the business operating model and strategic priorities. I ask, "What's important for the business? What's their strategy now and in the short and long term? How are technology and data enabling the business vision, contributing to strategic outcomes, and potentially hindering progress?" For instance, in my role in the cinema industry, I analyzed consumer data and used it to drive revenue.

I approach it with two lenses, examining what the business is doing with data and technology to both drive revenue and operate efficiently. My goal is to identify current challenges and potential opportunities that the company can capitalize on. Looking again at an example from my cinema experience, I looked at labor optimization and movie scheduling, using data to understand audience demographics and preferences. For example, you schedule more Bollywood movies in Birmingham than you would schedule in other areas due to the higher Indian population. Knowing the data around demographics in specific areas is invaluable there.

LISA: How did you overcome challenges and get businesses started on their tech and AI journey?

NADINE: My approach involves understanding the business model, vision, and priorities and mapping stakeholders' responsibilities. With this understanding, I take stakeholders on a journey, showing them how technology and AI could drive more revenue and impact the business. The method of driving this change depends on the organization and specific program.

NADINE THOMSON

LISA: Did you drive it through the business leaders or the technology function?

NADINE: It differs depending on the organization and the program. So, for example in the cinema industry, we brought in data and technology to optimize labor, an initiative that the business had already prioritized. As well as collecting the data we looked at some external factors too and brought them together with the scheduling of movies and labor.

We looked at what we needed to put in place for ticketless cinema, which competitors were starting to do (this was 7 or 8 years ago), and used that to drive efficiencies within the business and also to drive more sales because our research showed that people were more likely to go to a movie if they booked well in advance rather than deciding an hour before. Presenting this to the executive team helped drive the data and tech solutions.

LISA: Did you have the skill sets already in your team or did you have to go and recruit them or use external teams to help you?

NADINE: Again, this varies by organization and team. In my role with the cinema business, I partnered with external organizations to complement the existing capabilities, as the team might understand the business operations and data sources but lacked the skills to synthesize data or build specific solutions. Whereas when I have worked in media, they definitely had the skill set in their organizations.

LISA: What are the main obstacles you encounter?

NADINE: In the age of data abundance, I find that data quality and consolidation are among the biggest challenges. With storage solutions like the cloud making storage so easy, people now tend to collect everything. Businesses often have huge amounts of data but aren't very skilled at bringing it together cohesively, so one of the key challenges is consolidation. This is exacerbated when different

business areas label data differently. The issue extends from poor data quality due to human error during data entry to joining disparate data together. You can only realize the potential value of enormous amounts of data when you have a brilliant modeling system that can work across all that data and pull something meaningful together.

LISA: Could you share examples of how you've implemented a data product within a business and the impact it had?

NADINE: When I worked in media, we developed a product that analyzes client advertisements. It suggests tweaks, such as moving a logo from the top left to the bottom right, to potentially increase the attention the ad receives. It's an AI-backed system, drawing on past ad performance data and identifying patterns. However, the challenge is that AI suggestions don't always align with real-world considerations. For example, if the AI suggests including only a male figure in a soccer shoe ad because historically it attracts more clicks, we may not follow that due to other factors. So, while AI is an integral part of what we do, it doesn't take over entire processes or products.

LISA: What are your thoughts on the current state and future of AI?

NADINE: Although AI is commanding a great deal of attention now that it has a user interface with ChatGPT, it is still early in deployment within businesses. What people claimed AI could do years ago is just becoming feasible now as its extraordinary capabilities usually need a few years to become available at scale.

However, things are evolving rapidly, especially as AI solutions like ChatGPT have captured the imagination and are encouraging organizations to lean in, propelling AI's development.

NADINE THOMSON

LISA: What, in your opinion, are important considerations when working with AI?

NADINE: To me, the most crucial aspect is looking at human moderation. Regardless of how you're incorporating AI, whether as part of a tool or in a predictive model, it's important to apply human judgment, like that example I gave of the football boot ad. AI might provide suggestions based on ROI or past data, but they're not the only factors you make a decision on. AI lacks the ability to consider broader context or evolving cultural nuances.

For instance, in a project we conducted with a cinema chain, we used AI to model potential audiences for films based on past performances. The AI struggled when we introduced new elements, like a movie with a female or black lead. There was limited historical data for such cases, which led to the AI potentially underestimating their potential success. The AI couldn't predict their success because it didn't have past data to draw upon. While the AI wasn't biased in a malicious sense, it could not consider the data in a contextual way, taking into account the global, political, and cultural shifts that could influence the film's success. This is a limitation of AI that we must keep in mind. It underscores the importance of human moderation to supplement AI's outputs.

LISA: How long does it typically take to get started on a program such as cinema insights and generate something impactful for the business?

NADINE: It largely depends on the infrastructure already in place within the organization. If programs are underway, and significant work has been done to consolidate data across the organization, then things will move faster.

When introducing something new, particularly in AI, the quality of the existing data is crucial. A large portion of the effort is spent behind the scenes, cleaning data and organizing it in a way that allows us to implement a model and extract meaningful insights, like matching up audiences correctly in that media planning example I gave earlier.

In fact, I'd probably say that about 80% of the effort is dedicated to these seemingly mundane, yet essential admin and structuring-type tasks, while only 20% goes into implementing clever technology. There are some new technologies you can layer on top, which do some of this type of admin work for you, but it's not plug and play by any means.

LISA: What do you consider one of the most worthwhile inputs or interventions you've had during a data initiative?

NADINE: For me, the key is educating people to understand the significance of initial manual data entry and the subsequent joining of data to extract meaningful insights. It's crucial to recognize that data must come from somewhere and must be connected correctly to be useful. Getting that understood and bought into by an executive team early on is crucial otherwise they just don't understand why things take so long.

LISA: Can you share an instance of apparent failure that has equipped you to use data more effectively? Is there anything you would have done differently?

NADINE: I came into an organization a couple of years ago who were heavily involved in business intelligence reporting, but the business didn't understand why it took six months to deliver results. Some of that was due to the efficiency and structure of the team but they were also dealing with data coming in from so many different sources and in lots of different formats, which was a mess and consumed a lot of time.

The advent of AI and data modeling has made data matching easier. However, it is important to verify that AI has correctly matched the data to your business requirements. The executives in this case and many others think data is an instantaneous solution – put in data, hit a button, and a report pops out – when actually, there's about three months of work that needs to go on before you can actually hit that button.

NADINE THOMSON

Over the years, I've prioritized securing sponsorship to ensure you have the support to see you through the entire program. The executives need to clearly understand that a lot of groundwork happens before the report is generated.

LISA: What are you currently most excited about in AI?

NADINE: I'm excited by AI's potential to expedite processes by automating the more mundane tasks of data joining. Generative AI, for instance, is astounding. Recently, I was looking at a product that can take a picture and modify the background based on user commands, while avoiding copyright issues. This is something that would usually take a long time and, although it still requires some tweaking from a brilliant creative to pull it all together, the speed of change is both daunting and exhilarating.

LISA: Is there anything else you'd like to mention for business leaders?

NADINE: I'd advise board members to take more interest in data and AI. Boards predominantly recruit individuals with a finance background, which is understandable but lacks balance given the rapid rate of technological change. Boards should consider whether they have the necessary skills and support to guide them on data usage and AI strategy. They need to think not only about using data and tech but also about understanding it.

Key Takeaways

- **Vision-driven change for value:** Understand the business operating model and strategic priorities, and how data and technology can enable them. Have a clear vision of how data and technology can add value to the business, both in terms of generating revenue and improving efficiency.

- **Adaptive engagement strategies:** It's crucial to engage with stakeholders across the organization, showcasing how technology and AI can influence business outcomes with relevant examples. The approach should be flexible and adaptable, selecting the optimal method of instigating change based on the organization and specific program, and utilizing internal capabilities or external partnerships as necessary.
- **Data quality and consolidation challenges:** Data quality and consolidation are major challenges for businesses that have abundant data but lack the skills to integrate it coherently. Data quality issues can arise from human error, inconsistent labeling, or disparate sources, which can affect the accuracy and reliability of analysis and insights. AI enhances data products by optimizing ads, but it may overlook real-world factors like ethical or societal implications. While AI improves processes, it doesn't replace human judgment or creativity. A balanced approach combining AI and human expertise is essential.
- **Human moderation and AI limitations:** Human judgment is essential for evaluating AI suggestions, as AI may not account for broader context or evolving cultural nuances that affect human decisions and preferences. AI may have difficulty predicting the success of movies with female or black leads, due to limited historical data and lack of contextual awareness of the changing social and cultural landscape. AI is not biased or malicious, but it has a limitation that we must acknowledge and address by providing more diverse and representative data and feedback.
- **Data infrastructure and quality:** The time it takes to start a data program depends on the existing data infrastructure's quality and organization, which is vital for new data products, especially in AI. Much effort is spent on cleaning and organizing data to facilitate model implementation and extraction of meaningful insights. Data education and stakeholder buy-in are essential for getting people to understand the importance of accurate and consistent initial manual data entry and later data integration for value extraction. Securing early executive support for the data program is essential for aligning expectations on timelines and anticipated benefits.

- **AI on boards:** Data and AI are transforming the business world, creating new possibilities and challenges for organizations. Board members need to be more interested and involved in these areas, as they have strategic implications for the future of their organizations. However, many boards lack balance and diversity in their composition, as they tend to recruit individuals with a finance background. This may limit their ability to understand and guide the data usage and AI strategy of their organizations. Boards should consider whether they have the necessary skills and support to make informed decisions on these matters and seek external or internal expertise if needed.



10

Sarvarth Misra

Co-Founder and CEO of ContractPodAi

I've often been questioned about the logic of pursuing a law degree before creating a tech product. But after seven or eight years of doing this work, I realize the value of my legal domain experience. With AI technology becoming increasingly accessible, domain expertise, in many ways, will be the differentiator in the coming decade.

LinkedIn: sarvarth-misra-69869953

SARVARTH MISRA

Sarvarth Misra is a prominent thought leader, innovator, and entrepreneur in the legal sector. He is a Co-Founder and CEO of ContractPodAi, a market-leading legal technology business at the forefront of AI-driven solutions. Sarvarth was named a Champion of Generative AI Innovation by the Financial Times in 2023.

With a background as a corporate and commercial lawyer, he has extensive experience in the legal technology and services sector and has been instrumental in the development, strategy, and growth of ContractPodAi. Leading its design, development, and all new innovation initiatives, Sarvarth has also spearheaded its commercialization strategy and new market routes. His previous work in private practice involved working with leading law firms in Asia, focusing on M&A (Mergers and Acquisitions), capital markets, corporate restructuring transactions, and cross-border joint venture agreements in the insurance, IT, and financial sectors. An alumnus of the University of Oxford, Sarvarth is a prominent speaker at conferences on legal technology and services and has written on diverse topics such as AI adoption, the future of the legal profession, the impact of generative AI, company agility, and risk management.

LISA: What were the motivations behind the creation of ContractPodAi and what challenges were faced during its development and launch?

SARVARTH: With a background in commercial law as a corporate M&A lawyer, not technology, I've spent a significant proportion of my career reviewing, drafting, and negotiating contracts. My co-founder, Robert Glennie, has a similar background managing a global consortium of law firms for KPMG. Our experiences made it clear that the existing systems for contract management, which often focus on either buy-side or sell-side contracting, are not fit for purpose. This left the lawyers in the middle frustrated. We noticed that legal functions were buying technology to help manage contracts but were still dissatisfied. The legal system has to innovate in a modular way, so you have people buying different parts from one place, something for the repository, something for the workflow. Many of these companies weren't managing contracts effectively with next to no digitalization, which was surprising given

the size of some of the customer bases. Contracts were scattered across various repositories, even on someone's desktop. There was no single source of truth.

So, for instance, when lawyers started contract litigation, it was often the first time they were hunting through shared drives or emails to find contracts. The number of times people missed out on price escalations or renewals, which had significant top-line and bottom-line impacts on businesses, was substantial.

We thought, hang on, we're the guys building these contract management processes across enterprises that did not meet the needs of the business, so let's address that. Asking our own customer base as private practice lawyers, large GCs (general counsels) told us, "You guys understand what we do. You understand contracts. You know how the process works in an organization. You know how broken this market space is. Have you ever thought of building an end-to-end platform?"

Our view was that a significant majority of the market needed a solution to help them in their digital transformation journey. So, Robert Glennie and I put our heads together and launched ContractPodAi in late 2012, an end-to-end AI-based contract lifecycle management system to make contract lifecycle management fast, easy, and intelligent.

LISA: Who is your core customer?

SARVARTH: Our customers are every user in an enterprise who touches a contract in its lifecycle. That said, we pride ourselves in focusing on championing legal departments to drive transformation in contract lifecycle management. Our system covers everything from how the business requests for contract reviews or new contracts, to how you assemble and author your high-volume, low-complexity contracts. It also includes using artificial intelligence for third-party contract review and signature management.

SARVARTH MISRA

LISA: Can you give a brief overview of the services you provide?

SARVARTH: ContractPodAi offers a suite of features including Leah, a generative AI-powered legal assistant, and an intuitive Contract Lifecycle Management (CLM) platform for end-to-end contract management. It also provides robust analytics to optimize operations, a legal intake solution for request organization, and translation capabilities for over 60 languages. The company operates on a Software-as-a-Service (SaaS) model, offering continuous updates and support to its subscribers.

LISA: Can you describe your company's AI-led growth and business model distinction?

SARVARTH: As experienced UK commercial lawyers, this is where we started. Initially, we thought our market would be SMEs (subject-matter experts) or mid-cap businesses, as we assumed larger organizations would already have contract management systems in place. However, to our surprise, when we got the product to market, we started gaining traction with very large organizations like Bosch Siemens. It's the usual entrepreneurial journey – you think you know your market audience, but you end up somewhere else. The pain points were quite similar across the board – there was no single source of truth for contracts, the process was broken, and it wasn't done well.

Innovative GCs who were more tech-savvy than their predecessors found this frustrating and were eager for change. We began attracting these large companies with similar pain points, and thus, the product-market fit started veering into the late mid-market to enterprise space. Our first customers were effectively large businesses, which was quite accidental.

Fast forward to 2019, we established a presence in the US, specifically New York. Over the next 12 months, we saw similar acceptance of the product-market fit. We've taken a fairly unique journey in the enterprise B2B software space. This business hadn't taken any institutional money and was fairly bootstrapped. It had a reasonably good ARR (Annual Recurring Revenue) great retention

numbers and returning customers. This is quite different from most of our peers who probably spent \$20 million building a Minimum Viable Product (MVP).

Our company's growth was initially surprising to analysts due to our tangible revenue base with minimal funding. Despite the lack of a CTO and being founded by two lawyers, we managed to grow our product platform. We then partnered with global venture capital and private equity fund, Insight Partners to transition from a bootstrap business to a scale-up, focusing on growth with a scientific approach. We started building an operating model, considering sales, demand generation, enablement, and building a marketing organization from scratch.

From late 2019 to 2021, we maintained a significant growth rate and market traction. Then, in September 2021, we secured a substantial funding round of over \$100 million from SoftBank. This significant investment was strategically aimed at shaping the company's trajectory over the next five years and how we would devise a strategy from both a product and go-to-market motion to lead the legal technology market in our areas of strength. This funding raise allowed us to continue growing and innovating on the product side so we could serve our customer base with truly cutting-edge technology.

Our focus has been on enterprise customers, positioning ourselves as a legal lead contract lifecycle management player. We empower legal operators and counsel as decision-makers in buying enterprise technology that benefits them.

We've been able to maintain a good balance sheet. We have always been committed to growth that's underpinned by scientific principles, and this has consistently been our core focus. We've always wanted to make sure we become a little bit more efficient, and that's what's helped us along the way. We've always thought about sales, demand generation, enablement, and building a marketing organization from the ground up.

SARVARTH MISRA

LISA: Could you share your personal journey in AI and the development of this platform, especially from the perspective of a domain expert rather than technology?

SARVARTH: We were pioneers in contract management, even when it wasn't a trend, and embraced the term AI in our brand name in 2015. There are two main aspects to our approach. First, the power of workflow automation, which streamlines the process. Second, the use of artificial intelligence to eliminate manual data entry from contract management. This allows us to capture a wealth of data, which can then be transformed into valuable insights.

We recognized the potential of these technologies early on, and in 2015, we were among the first to partner with IBM Watson. We initially focused on three use cases.

We leveraged AI for tasks like data entry, which were traditionally done manually by non-lawyers.

From there, we delved into the concept of cognitive deep search, especially in the context of contracts. For instance, if you're a lawyer going through an M&A, you have thousands of contracts and you need a way to search for specific provisions across your repository. This became a major use case for us.

Additionally, when it comes to contract review, governance, risk, and compliance, there's a need to review third-party contracts consistently, which has been a gap in the market. This was another significant use case for us. We were probably ahead of our time in doing this, but as technology improved, the market caught up.

LISA: What are your observations on the expectations of customers for results?

SARVARTH: Our journey with these use cases and applied AI solutions over the past five years has had a profound impact. We were using natural language processing and machine learning, which are great, but legal documents are complex and unstructured, so showing value to customers typically takes time. It's a journey, not an overnight sensation. The increasing speed of AI

innovations, including ChatGPT Large Language Models (LLMs), has won many fans but also created an impatient mentality: if it doesn't yield something immediately, it's not worth anything. Sure, there are shortcuts and progress on the way, but AI needs time and investment to turn it into something viable for an industry. Customers who understand and accept that this transformation is always going to be a journey have aligned their expectations. However, some customers, who are not used to the crawl, walk, run journey, expect immediate results and lose patience.

So, our experience with AI has been the sum of seven years of successes and failures. Today, the LLMs are constantly improving. Add to that, our subject market expertise and deep engineering skills allow us to solve complex legal use cases – and that gets us excited about what the future holds.

LISA: What technology did you use to build your platform?

SARVARTH: Our technology stack was built on Microsoft Azure. This is because many of our products, including features like redlining, are built as plugins for Word. As lawyers ourselves, we knew how hard it is to get lawyers out of Word, so we decided to develop within the Microsoft ecosystem. This was a tough journey, but it eventually paid off as you get more stickiness because the lawyers will actually start using your software in many ways.

From the outset, in 2014, we were cloud-first. Our model is based on software as a service (SaaS), with a license subscription fee model. We've fully integrated signature partners, like DocuSign, into our workflow. But the core build, workflow, and AI are all our own. This blend of Microsoft tech and our proprietary technology has been instrumental in our success.

LISA: Could you clarify how IBM's Watson was integrated into your tech architecture?

SARVARTH: In 2015, we built our foundational model for contract management with Watson. We then trained our own customized model in much more detail with our customers' feedback and our own data science team. This investment in the data science team was important. Of course, our bandwidth of technology partners has expanded to Google, Microsoft, OpenAI, and Anthropic to name a few.

LISA: Did you find it challenging to recruit data scientists?

SARVARTH: A good data scientist is always hard to recruit. However, we found our edge in the ability to attract what we term legal engineers – individuals with a strong legal foundation who also possess a deep understanding of technology.

Legal engineers and legal designers working in tandem is now a common sight, as the need for subject matter expertise is undeniable. It was this combination that paved the way for our success in building our own models and more, as we went through this journey. Contractual language is a bit complex to figure out so you need expertise in many ways.

LISA: How did you ensure the protection and ring fencing of intellectual property? I can envision that the legal department of one enterprise would not want its contracts to be confused with those of another. How was this issue of containment addressed and resolved and incorporated into your architecture?

SARVARTH: Given our customer base, from Day 1 we differentiated between what we call embedded AI versus purely generative AI. With the embedded AI techniques, AI is being used to only focus on the embedded data of the customer versus publicly available internet data. So AI is being used to understand each individual customer's data in a secure environment and using engineering techniques like RAG, reinforced learning and prompt engineering to build customer-centric models. Our infrastructure layer, powered by dedicated compute resources, ensures that no customer data is shared for training public models.

Over time, trends from our customer-centric models became learnings that we could incorporate into our own proprietary models.

LISA: What has been the reaction of lawyers to generative AI, I suspect they were both excited and terrified about it?

SARVARTH: You're absolutely right. My peers urged me to explore GPT in its very early days and I very quickly realized that this is different. But I also realized that these models by themselves will not be able to solve legal use cases – deep domain expertise and deep engineering will be needed to show results at the enterprise level.

We knew this would have a transformative impact and that we had to adjust our product roadmap. We doubled our efforts over the next few weeks. Fortunately, our previous eight years of AI and enterprise experience gave us an edge. It was clear to me that language models would have scalability and therefore deployment was easier.

When we saw our competitors using it, we knew we were ahead. We had already tackled the challenges they were just beginning to face. One of the things we were very clear about from day one was to be proactive adopters of AI, not reactive. The decision I instinctively made for the company was to take the lead in integrating GPT as an AI-led company, understanding that one size wouldn't fit all given the multiple legal use cases.

We knew integrating a GPT model wouldn't be a quick fix. It would require deep engineering and be painful, but it would be worth it. Recognizing that lawyers were initially experimenting with GPT from a consumer, non-legal professional perspective, we understood that there would be skepticism and concern regarding risks. We needed to do this integration in the right way.

Legal councils don't want Google or Bing answers; they want insights from their own data. So, we embedded these models into our system and quickly realized that OpenAI's standalone API wasn't equipped for this. It couldn't handle large volumes of data. So, we undertook a massive engineering effort, including

SARVARTH MISRA

our own dedicated compute for processing, to optimize output performance and security for our clients. We knew we needed this infrastructure, despite the investment. In hindsight, moving with conviction was one of the better decisions I've made.

LISA: How did you work with OpenAI and what impact did this have on infrastructure costs?

SARVARTH: We worked directly with OpenAI. The dedicated compute was pricey, but it was a necessary step for serving enterprise customers because we didn't want our customers to be just another public API. That model might work for B2C or consumer apps, but not for our clientele. Early on, we recognized that infrastructure costs would accumulate over time. But that's a part of the journey we've embarked on. It's all about the long game.

LISA: Have you maintained your partnership with OpenAI, or are you considering collaborations with Gemini or other conversational AI models?

SARVARTH: We've multiple AI strategies, including with OpenAI, Anthropic and Google. We've also tested some very good open source models that we have been able to train for different legal use cases. Given our experience with AI in the past several years, our strategy was very clear that "one size doesn't fit all" – even with the different complexity of legal use cases, we will need to use different models and engineering techniques to deliver the best end result to our customers. Despite the strengths of these models, we realized the need for a multiple-language model strategy to cater to the depth and breadth of our solutions.

We made a larger investment in enterprise cloud infrastructure than most, which we believe was the right thing to do. This in many ways signified our moment to lead, aligning with the business's purpose and mission focused on AI.

SARVARTH MISRA

LISA: How did you convince lawyers of the merits of GenAI and integrate it into your platform?

SARVARTH: I've trained as a lawyer, so I understand the legal persona of being risk averse. Skepticism has always been a challenge but seeing their peers embrace change and discussions with CEOs and boards pushing for operational excellence made a difference. We leveraged our experience in AI to give an honest view to our customers about what can and can't be done. Managing expectations is very important and we continue to invest in educating people about the possibilities and limitations.

When the CEOs of enterprises started to prioritize GenAI and started asking their GCs how they would deploy it in the next 12 months, this shifted the momentum and marked a fundamental change in priorities. We were getting approached by GCs saying, "My CEO wants me to use GenAI."

To integrate it successfully into our platform, we've considered how GCs might use AI, given that many CEOs are keen on adopting it but often without a clear application in mind. We've modularized our day-to-day operations, building modules for tasks like drafting and redlining. Transactional work and guidance are key examples. We estimate that for most GCs, these tasks constitute 70% of their daily work. We believe in leading with a human touch, as we're not yet at a stage where we can fully rely on AI as an autopilot technology. We've strategically positioned our AI to handle tasks that lawyers typically don't enjoy, like answering routine business queries or reading contracts. If some of this work can be handled more efficiently with the help of technology, it could lead to significant savings. This is our vision for the future of legal work with AI.

At the end of last year, we won an important industry deal with Integreon, a provider of outsourced legal and business solutions. They will use our GenAI legal platform Leah to transform and reimagine the delivery model of some of their own legal service delivery lines. We have signed a similar global enterprise deal with a leading Big 4 consulting firm as part of their legal practice. Both these deals are highly impactful examples as companies look at reimaging their operating model for legal departments with a GenAI-led strategy.

LISA: What advice would you give to leaders from non-engineering backgrounds on how to enhance their knowledge and skills in technology and AI?

SARVARTH: I think vertical expertise matters as technology becomes more democratized, where models will become much more readily available, and coding will become easier. I think when you are solving problems, you need to be an honest student of the subject at hand. Thinking deeply about problems within any domain is essential. I've often been questioned about the logic of pursuing a law degree before creating a tech product. But after seven or eight years of doing this work, I realize the value of my legal domain experience. With AI technology becoming increasingly accessible, domain expertise, in many ways, will be the differentiator in the coming decade. Understanding contracts and the legal system has been invaluable. I understood what could be commoditized.

I would not underestimate the potential of individuals with diverse backgrounds in biology, healthcare, finance, or any other field to excel in technology. Domain expertise is sought after, and with technology becoming increasingly accessible, it will be the differentiator in many ways in the coming decade.

While general management skills are important, subject matter expertise is equally crucial. As I look across vertical specialist leaders today, they have a mission that goes beyond just starting a company. They have a story, an experience that informs their approach. This experience, even if it's from 15 years ago as a lawyer, comes back to you and helps you understand your buyer, who in our case is a lawyer.

Lastly, the other thing I always believe in is the importance of core competency as a strategy. I think it's easy to sometimes diversify a business too fast. But I think focus, expertise, and competency are crucial.

LISA: What impact do you foresee in the next five years?

SARVARTH: The adoption of generative AI is not just about staying relevant, but about leading the charge toward a more efficient, value-driven future. It's about reimagining the way we do business, the way we serve our clients, and the way we shape the market.

The change management that will be driven by generative AI adoption across our customers is most definitely once in a generation. It will need customers, technology providers and consulting firms coming together to reimagine the future state of an enterprise and its operating and delivery model. Investments in technology infrastructure, responsible usage, and compliance will also increase. When I then think about SaaS providers, we will be pushed to the cutting edge of technology but also will need to demonstrate more than ever before the value we bring to our customers to stay relevant and thrive.

Key Takeaways

- **Long-term investment:** AI can be powerful for automating manual tasks, conducting deep contract searches, and ensuring consistent reviews of third-party contracts for governance, risk, and compliance. There can be quick wins along the way but the implementation of AI into core processes is not an overnight sensation. Patience is required to realize its full potential and value.
- **User-centric design:** Understand the workflow of your users as well as their technology preferences and build these into your solution using existing technology where it makes sense to increase engagement.
- **The role of domain specialist engineers:** Hiring specialist engineers – those with domain and tech expertise – to work in collaboration with domain process designers, leads to successful proprietary model building.

- **Embedded AI for secure-centric models:** “Embedded AI” uses techniques that keep customer data secure while leveraging it to address specific customer needs. Engineering methods like RAG, RL, and prompt engineering personalize the experience for each user, and dedicated computing resources ensure this data remains private and isn’t used to train external models. Over time, general models can be improved by incorporating trends identified in customer-specific models within a structured system that maintains data separation. This approach allows for continuous improvement based on customer feedback and trends while safeguarding the unique nature of customer-centric models.
- **Targeted AI application:** Focus on high-value processes and modularize the underpinning operations by identifying specific, routine tasks within your industry and then select the most appropriate AI models for those tasks.
- **Domain expertise:** As AI becomes more accessible, domain knowledge will be a key differentiator. Leaders from non-engineering fields should value domain expertise. The potential of individuals from diverse backgrounds to excel in technology shouldn’t be underestimated. While management skills matter, subject expertise is equally important. Core competency remains a crucial strategy.



11

Edward Fine

**AI and Data Science Consultant, Technologist,
and Instructor**

The lovely thing about data is that there are always younger people, those who are earlier in their career, or those who are just data-excited who have got the bug and want to see what data can do for their companies. That natural excitement is the best ally that you have.

LinkedIn: [edward-fine-239a28](#)

EDWARD FINE

Edward Fine is an AI and data science consultant who develops technical solutions to real-world problems. He is currently an AI instructor at Afinepoint, but prior to this he worked in fintech, eCommerce and academia, primarily as the Faculty Director of Data Science Essentials at UC Berkeley. He specializes in developing creative, actionable, analytical insights in results-driven environments.

LISA: How did you get into data science?

EDWARD: I started out in quantitative finance back in '98 and did a lot of statistics and optimization through the financial crisis. By 2008, I realized I was losing interest in finance but found the data scene in Silicon Valley enticing. Knowing some of the earliest data scientists helped me move to this area, bringing the same skills to different verticals – it's been a constant theme throughout my career.

LISA: What is your role now?

EDWARD: I've worked extensively in fintech, e-commerce, healthcare, education, and natural language processing. Having taught various university courses, mainly at Berkeley, I now serve on the external advisory board at San Francisco State University's computer science department. Obviously, Berkeley is a top-tier university, but Cal State is for everybody, so it has a very diverse student body with a fun set of challenges, which I really enjoy. I also evaluate AI companies.

LISA: What examples can you share of how you convinced companies to prioritize data?

EDWARD: I was brought in as the first lead data scientist at a fashion startup called Kings Lane. It was a very brand-driven company, so you could make a strong case for data but trying to get artistic, fashion types to buy into it was a challenge. So, my strategy was to introduce a low-stakes, internal data product, focusing on cost savings rather than revenue enhancements to build trust.

EDWARD FINE

Of course, the bottom line is never as “sexy” as the top, but with a relatively simple product, we provided analytics on customer purchase patterns to the inventory buying team. This was a flash sale type model where the buyers needed to know what actually sells. Because it was an online store, we had a SQL database with what had actually sold, but it wasn’t accessible to the back-end buyers. We started with rugs and asked which tended to sell, 2m x 1m or 3m x 5m? Then I was able to make an initial win with a sequel query that would run that, then I put it into a product and gave it to the buyers so that they were able to look at what had sold and how consistent a bestseller it was. Standard deviation turned out to be tremendously powerful to them because these things would always sell, but did they have to put it on the back tier where there was a greater markdown? What products would sell on the first sale with the highest selling price and what would end up selling later on? We were able to have them alter their buying plans and that got us fans on that side of the house.

LISA: How specific could you be about what sold well?

EDWARD: It was interesting to see where the strengths and limits were with this. Because this brand was in a very fast-changing industry, historical data analytics on things like color were not predictive because trends were always changing. It was a very interesting piece of wisdom because the types of insights we data guys provide are great in stable environments, like those you’d find in hardware screws, for example, but something as trend-driven as home décor is not as stable an environment – last year’s average isn’t going to tell you much about this year’s. So, as a data person in retail and especially in e-commerce, you need to know what things will change and what will stay the same from this year to the next. While you can predict seasonal sales for things like umbrellas in the rainy season and gloves in winter, fashion trends require a level of respect for the experts whose job it is to stay on top of them.

EDWARD FINE

LISA: What's your approach to solving problems in a company?

EDWARD: I aim to identify a specific mandate with a specific group. If you find yourself in the unfortunate situation of hearing, "Here's our database, add value," that's tough because it lacks specificity. So, your job is to strike up casual conversations (real or virtual) with individual contributors or low-level managers, to understand their pain points and identify ones you can alleviate. You aim to build credibility with individuals, but the ultimate goal is to demonstrate value quickly by addressing a problem that significantly helps the business. It's a tricky sizing piece, and one where junior data people often get tripped up because you don't want to do something that's too big because it will meet with internal resistance, but you don't want to do something so small that it doesn't make a difference.

Ideally, if you can pay for your first six months of salary within a couple of weeks, that's a great win.

LISA: Can you explain how you build a product and drive internal adoption?

EDWARD: The first thing I did was run a query, take a snapshot, and drop it into a report, pushing it out in an automated email to business stakeholders. This was how I did it back in the Dark Ages, and because of our on-premises security model it was not easy to build a website around it. This was probably around the time Yammer was purchased by Microsoft, in 2012. That was good enough to deliver a Monday morning report. The data actually changed, probably once every six weeks, so an email format with this slow-moving data was fine in those days. What I would do is take on some kind of driven analytics system – there's a whole bunch from Microsoft – or I would build it out with a fast API these days. Now I would probably just go to my favorite GPT and tell it what tools I wanted to use and have it built out as a front-end website for me.

EDWARD FINE

LISA: What mistakes have you made and lessons learned?

EDWARD: Oh, so many. I've made a career of it! Patience, humility, and strategic positioning of quick wins in adjacent areas can build trust and gradually steer the organization toward a data-focused culture.

Each new assignment or role in a company requires understanding its unique corporate culture and recognizing resistance points. Any attempt to bring in data-driven changes should respect these dynamics. Finding a few key allies at the start helps you assess what's really important and understanding the culture will give you the context of whether the changes you bring will be accepted. If the culture is that way for no good reason, then you can probably work with it. But if it's that way because they once had a million-dollar screw-up, then you're going to have a tougher time.

LISA: How can a senior business leader drive change amidst the intimidating jargon?

EDWARD: Encourage curiosity and openness in understanding new terms and concepts.

The secret advantage everybody has is their young people. Leveraging the energy and drive of young, data-excited individuals in the company is crucial because they'll put in the extra work to explore what data can do in their companies. They see that it brings a natural advantage and helps them become a more valuable employee, in their current company and in their future career.

And, when it comes to feeling intimidated, the biggest obstacle is that people don't want what they don't know to be exposed, so I lead by example. One of the things I learned at Microsoft was to be ruthless about asking what three-letter acronyms meant. The first time you do it, you might look silly. But afterwards, you get known as somebody who, if they don't know something, is willing and curious to ask.

EDWARD FINE

LISA: What excites you most about the future of data?

EDWARD: I'm thrilled about the potential of generative AI and its capacity to process specific training data. Right now, we're seeing the amazing power of generic large language models, but the power of applying those models, fine-tuning them, combining them with an individual, or an enterprise, a company, a startup... that data is going to be amazing. Imagine a world where we can capture data from our surroundings using our phones and have AI process it and provide valuable insights. Say, I walk around my garden with my phone and AI gives me advice about whether I'm watering too much or whether the citrus trees need more iron. Or maybe I can take a picture of that thing that's not working in my home and it will recommend the right video to show me how to fix it or customize a video to my specific skills. I can't wait for that!

LISA: What's your advice for building an effective team?

EDWARD: It's rare to find a full stack data scientist. They do exist but they tend to be consultants and very expensive, so aim to build a team with a mix of business, mathematical, and engineering skills, the latter in someone who knows the architecture of data in the company. Borrowing engineering talent from other teams can work initially. You can then focus on enhancing their skills if they show enthusiasm for data.

LISA: Do you see the skill sets people need changing because of AI?

EDWARD: Yes, and it depends on what stage you're at. There's a need for skilled engineers who can help with the integration of complex data systems, like a data ECU, Pekin AI, or other types of data analytics as a service product. These are becoming increasingly important in the business domain. A new business professional will often find that these tools either have not yet been established or, if they have, they are concentrated in one specific area; for example, Google Analytics or similar web analytics tools monitor the front-end and are very good at that. Or, a CFO might have sophisticated analytics tools on the financial

end, but they are often specialized and restricted to their particular vertical. So, the changing landscape necessitates a range of skills – from understanding and implementing AI and data analytics tools, to integrating them across different verticals for a more holistic business approach.

Key Takeaways

- **Careful targeting of use case:** Swiftly demonstrate value by solving a significant business problem. Ensure the use case is strategically sized – too large may face resistance, while too small may lack impact. This is a common pitfall for junior data and AI professionals.
- **Building effective teams:** AI is reshaping skill set needs, increasing the demand for engineers who can handle complex data systems. Teams therefore need a mix of business, math, and engineering skills. Due to the scarcity of full stack data scientists, it's beneficial to borrow and nurture talent from other teams. Developing the ability to understand, implement, and integrate AI and data analytics tools across various verticals delivers a comprehensive business approach.
- **Building trust through low-stakes data products:** Building trust with data starts small, such as introducing a low-stakes, internal data product focused on cost savings. By understanding the pain points of a specific group within the business and addressing them with a data-led solution that quickly demonstrates value will establish momentum. Patience, humility, and focusing on these initial quick wins in adjacent areas will lay the groundwork for a data-driven culture as a prerequisite for AI solution development.
- **Navigating corporate culture:** Understanding the unique corporate culture in a company is crucial, as is recognizing resistance points. Finding a few key allies at the start helps assess what's really important, and understanding the culture provides context for whether the changes will be accepted.

EDWARD FINE

- **Role of senior leaders:** Senior business leaders can drive change by encouraging curiosity and openness. Create an environment where it's okay to not know everything and to learn together. Understanding new terms and concepts is crucial. Lead by example, such as asking about unfamiliar terms.
- **Harness the energy of younger colleagues:** Recognize the natural enthusiasm of younger individuals, those at the early stages of their careers, or those simply excited about AI. Their willingness to explore its potential not only enhances their value in the current company but also bolsters their future career prospects.



12

Sanjeevan Bala

Group Chief Data and AI Officer at ITV

We developed the last mile approach to shorten the time to value by focusing on the last stage changes, including business processes and functional needs. We made the outcome value the last mile and worked backward, instead of what you often see, which is people collecting lots of data, then building a lot of technology, then putting in data scientists.

LinkedIn: [sanjeevanbala](#)

SANJEEVAN BALA

Sanjeevan Bala is the Group Chief Data and AI Officer at ITV. Prior to this he was Head of Data Science at Channel 4 and has previously held roles within PwC, Dunnhumby, and a number of startups in Silicon Valley. He is also a Non-Executive Director of a Multi-Academy Trust where his role influences how we shape education to secure the next generation of AI talent in schools. He is also a FTSE 250 NED of a global B2B food business.

LISA: Could you briefly describe your career up to now?

SANJEEVAN: My career falls into four phases. The first was classic customer-centric management consultancy, asking, how do you put the customer at the heart of the organization and what does that mean for the operating model and customer journeys? Then I went into launching digital businesses during the dot-com bubble, working with entrepreneurs to get funding and then building digital businesses and scale-ups. For example, at Citibank we launched their direct consumer City X product proposition model in the US.

Next, I shifted client side to Channel 4 and eBay. As I'd done consultancy work, I knew what I was looking for and how to apply all these learnings internally. The big focus was on culture change, executive mind shift, and delivery of the strategy.

The last phase of my career has been in a non-exec role, working out how to amplify my contribution at board level, and ensure skills and understandings lead to value creation and, critically, prevent value loss.

Now, I'm the Group Chief Data and AI Officer at ITV driving growth through data and AI across studio and regular TV business segments. ITV produces content, promotes it, distributes it, and then monetizes it. My role is finding value cases across that entire value chain where we can use data and AI to drive growth.

LISA: How do you prioritize where to use data to solve business problems and how do you get that backed internally?

SANJEEVAN: At ITV, we look at the value chain, evaluating core activities such as content production, promotion, distribution and monetization. For each of these processes, we ask, “How might data change increase or protect value?” For example, if we could build a predictive model, how might we use that to acquire customers more effectively? And, if we could do that, what sort of lift might we see in marketing? We generate a lot of value cases across the business working with different stakeholders ensuring buy-in from all parts of the business. The accountability of these value cases remains with each of the functional leaders.

We start with individual use cases to understand their commercial viability, co-signing the value potential with each business leader. Next, we consider augmenting the existing revenue forecast, for example, where can you use data to bring in viewers at a lower cost per acquisition? Finally, we look at the net new revenue beyond the original five-year plan. You need to take a commercial and value lens to the CFO and convert those numbers to aligned objectives in the annual revenue plans and targets. That’s how we got alignment in the organization.

LISA: How did you organize the data team and costs to achieve this?

SANJEEVAN: Firstly, we decentralized data skills, embedding data scientists, engineers, and analysts into business units. This was key as it fostered ownership of the functions and the responsibility for delivering value wasn’t with the data team but their business unit. Integrating data into the teams so they all act together also built trust. They saw the data team as part of their team.

We’re currently moving costs to the business units to activate ownership of the investment and the value realization. This overcomes the unnecessary tension in the organization where the data unit is promising all this value but has no control of the levers. You need the ownership and the cultural alignment where the costs are being incurred, but also where the value is being realized.

LISA: Do you have communities of practice in the center?

SANJEEVAN: Yes, the center acts as the glue between different domains and different organizational components. For example, with projects initiated in marketing, we can quickly take a commercial overview and say, “If we change this data product slightly, it can have repeatable applications.” We’ve effectively taken the “data as a product” route, which lets us identify and introduce the best standardizing practices. Rather than building a data lake, we’re introducing a data mesh architecture that decentralizes control and everything into the business unit, enabling the business to own technology and data responsibilities. This type of architecture enables the business units to control the following: how shared data is accessed, who accesses it and in what formats.

LISA: What was your starting point at ITV?

SANJEEVAN: I started at ITV in 2020 just as lockdown started and the whole strategy and board engagement was done virtually and remotely. When I joined, the business understood data was important on a philosophical level, but it was quite nascent and foundational. They had a classic starting point of, “Let’s collect data because we sort of feel it’s important and let’s put it somewhere.” This meant that there were various degrees of data at different quality levels. There was registration information and data on what people were viewing on what was then ITV Hub. This was physically made available and put into a cloud environment. They had some resources, like an analyst and a couple of data scientists. The bit that was missing was the explicit link to the TV strategy, which then and now is around three pillars – supercharging, streaming, and monetizing our studio IP. We linked these pillars to the data strategy which helped us try explicit use cases and value cases. We prioritized the value case over the use case because it compelled the business and the data unit to rethink what it would take to fully realize this value.

LISA: How did the “last mile” approach develop and get adopted?

SANJEEVAN: We developed the “last mile” approach to shorten the time to value by focusing on the last stage changes including business processes and functional needs. We made the outcome value the “last mile” and worked backward, instead of what you often see, which is people collecting lots of data, then building a lot of technology, then putting in data scientists.

We proved the viability by taking a “thin slice” approach, building vertically instead of horizontally – focusing on specific value pieces. We asked, “OK, for this bit of value, what is the minimum amount of data infrastructure model that we need to go end to end?” So, we’d collect the data, tag it up on our website, build a model, put it through a security solution and email our website users to see if we can get them to activate. Once we prove that, we enable iOS, then Xbox, then PS4, and iteratively do it like that.

This approach meant we were hitting “last mile” challenges early, which meant we could pivot and iterate the strategy. We learned lots of things vertically in the stack, rather than horizontally, for example, problems relating to data quality, connectivity into the matrix pack, about the difference between the universe population and the available population. Drawing on our understanding of connectivity, we could integrate what we learned into the design and functionality of social platforms. Also, our insights into attribution could be leveraged to enhance social platforms and user interaction analysis.

Once we’ve promoted on TikTok, how do we measure whether they came back in and watched our content? These “thin slice” vertical learnings are how we got traction and buy-in from the business.

LISA: What was the shift in how data programs were organized before and after you joined the business?

SANJEEVAN: Before I arrived, the programs were tech-driven because a lot of it was cloud compute and distribute, so there was a lack of connection to business strategy. Once the business understood the connection to the corporate strategy and revenue and margin implications, it became a much more business-driven initiative.

My unit sits in the business, not in technology so we developed a trio strategy between myself, the CTO, and the business leaders. We all had slightly different objectives which was key. The business was responsible for the “last mile” adoption, organizational change, business change and realization of value. The data unit was responsible for the art of the possible, the “what if...?” and introducing a new narrative. The technology teams are responsible for reducing time to business value. This shift led us to partner and buy capabilities rather than build everything, which allows us to maintain modularity for future needs.

The technology team really shifted their thinking about how to reduce time to value rather than build. A lot of businesses’ functional problems have now been commoditized by a technology solution, so you’ve got to be really clear on the business case if you are going to build something internally. We have adopted a classic LEGO strategy – a modular approach to everything we do – so we can pull one bit out and put something else in down the line.

We helped core engineering teams understand higher-order assembling versus building. This was an interesting process because as a lot of the Software as a Service (SaaS) providers’ spaces are being rapidly commoditized, they’re going further up value streams and adding on capabilities to protect their margins. To ensure modularity, we were careful about the core capability we were buying; for example, we’re seeing SaaS providers bundling the capability with the tracking which can be really hard to unbundle in three years’ time if you are reliant on the tracking.

LISA: How did you ensure thinking around planning for horizon two (emerging opportunities) and three (future growth ideas) technology, while managing horizon one (core businesses providing most profits and cash flow, with a focus on improving performance)?

SANJEEVAN: We have a design authority community who do a lot of horizon scanning through the Digital Advisory Council (including external CEOs and non-executive directors), to look at what other technologies and trends are coming up and to ensure when engineering and architecting that we preserve modularity. The result is an injection of progressive, knowledgeable thinking, about trends like blockchain, Web3, and generative AI. We don't always get it right but we're always learning.

We've also navigated challenges in pricing with SaaS providers because they're scaled on a usage model. It disincentivizes us as a business to use it at scale, which has led to strange internal conversations. For example, we'll sign up to a recommendation engine, which has higher pricing the more we use, yet we find ourselves saying, "How do we then reduce the usage?" But really you want to put it everywhere to maximize the value.

LISA: How do you prioritize AI use cases?

SANJEEVAN: At the moment, we view AI as a solution to a business problem, focusing on the value and business outcome. We're currently in the experimentation phase looking at how it might augment productivity and the creative process. There's been an interesting shift for us regarding generative AI, which we describe as a distinct sort of subset, because it's suddenly allowing us to create value without having to be reliant on foundational high-level quality data. For example, what set of skills are required to think about the next drama plot, the storyboard, and the narrative arc, and what talent could be appropriate for these things? Then there are the potential uses in post-processing like special effects, additional edits, multiple language translation, and audio lip-synching, which we've already done a lot of trials on and are using extensively.

How it might be used in the creative process is currently less developed but there's certainly an appetite for how it might be used, as a co-pilot tool – it still needs a creative leader to drive it in order to create useful outputs. We're experimenting to understand if it gives us different ideas or allows us to do ideation faster.

LISA: How did you build up a team over time and retain them and were you dependent on third parties?

SANJEEVAN: To build capacity and get going quickly, we were initially reliant on third parties to build teams with diverse skills like data engineering and data scientists. The team model had a data product leader from a consulting background, who was business-oriented in terms of how they think and operate. Then we had a combination of agile delivery managers with delivery skills as well as classical digital skills – data engineering and data scientists. So that's how we built them.

We had a whole platforming team and had to assemble engineering and pipeline team analysts and data scientists, and then the whole data governance part, that's how we started building out. We focused first on three business areas – marketing, commercial and product. We did this because, with marketing, we can get the viewer in and get them to watch more content with commercials. With that additional content, we can monetize them at a higher yield, so we built the teams underneath that, against those particular value cases.

LISA: What advice would you give to a business starting to work with AI?

SANJEEVAN: I'd say, firstly, focus on the value case and work out what needs reporting, otherwise, you'll incur a lot of costs. Ask what you are going to do with it once you've done your reporting.

The second thing is to tell a story to help your functional leaders reimagine the “after,” as well as quantify that for your board and CFO to convey the commercial gain.

The third thing is the importance of horizon scanning because something you put in today could be commoditized tomorrow, so you need to avoid building where you don’t need to.

I would also emphasize the importance of change management, especially for traditional businesses. Change is 70%, while 20% is about the data and the kind of teams you need to build, and 10% is the algorithms and model. Often too much of the focus is on the AI without enough focus on the change.

It’s also key to solve the last mile first and sell it in; otherwise, you won’t be able to fully realize the value of all your investment. Joint ownership is important too, so, translate those lift numbers into personal objectives because that drives the right organizational behaviors.

Empathizing with stakeholders is also key, especially if an organization has low maturity with low churn, as starting with data/AI can be intimidating and daunting. So, walk in their shoes before you put your functional hat on around what’s possible.

My last tip would be to manage expectations. Clearly communicate your priorities and what they will and won’t deliver, as well as articulate the likely challenges.

LISA: Did you find any useful templates for value assessments?

SANJEEVAN: We developed an approach from experience and having been burned many times. Some methodologies exist but there’s no single playbook for this journey. Existing templates can be found, for example, we used a well-established one in our change methodology, but each organization must adapt them to their unique situation.

LISA: What have been the essential aspects of telling ITV's business story?

SANJEEVAN: There are three key artifacts. The first is how we define the value and revenue in terms of increment, protection and capacity. For example, "It's OK to get funding for a project that's in revenue protection, because if we don't do that, revenue values are going to go to somebody else."

Second, we use a learning cycle to foster organizational agility through iterative experimentation – starting with a hypothesis, assembling an experiment, running the experiment, capturing learnings, and feeding that back into the next cycle. We used a version of this visual cycle in every single functional part of our organization. For example, we'll talk about "If we can get more viewers in at a lower cost, and we predict who's gonna churn, we can target them on social media, we can bring them back in and we tell the story through that." That's been a very powerful visual.

Finally, the concept of the "last mile" has been essential in our change process. Now you hear the business talk about how we've got to fix the "last mile" first, which is an amazing change of perspective.

LISA: What are you most excited about now?

SANJEEVAN: The implications of what Web3 might do for data and digital-driven businesses are really fascinating. When organizations have collected data and used it to drive commercial outcomes or value outcomes, how does that value exchange change when you think about token economics in a Web3 construct and how do you think about creating value in that sphere?

I'm also really excited about collaborating across industries and combining different functional areas of innovation to open up new growth, for example, a media company with a retailer, or a media company with a financial services company. When you start to think about combining in terms of the underlying understanding of the consumer, then that creates a new order and unlocks new propositions that can be built on top.

LISA: Can you share any other key lessons learned?

SANJEEVAN: We overestimated adoption and underestimated the change challenge, especially in the “messy middle” as I call it – between the Gen Z digital natives starting at our organization and the board.

LISA: What would you have done differently?

SANJEEVAN: I think we would have embedded individuals earlier and been clearer about the consequences of prioritizing in specific areas. Also, I would have showcased more of how other organizations with similar functional processes are doing it, for instance, have a peer marketer talking to a marketer, rather than a data professional articulating something to a function leader. I think this would have made the conversation more accessible from the start.

Key Takeaways

- **Understand commercial viability:** Evaluate the value chain, including content production, promotion, distribution, and monetization, to generate value cases across the business. Begin with individual use cases to understand their commercial viability and co-sign the value potential with each business leader. Emphasize commercial viability in decision-making processes, view AI as a solution to a business problem, and experiment with AI to augment productivity and the creative process.
- **Augment existing revenue forecast and look at net new revenue:** Enhance existing revenue forecasts and explore new revenue generation avenues beyond the original five-year plan by leveraging data. Optimize current revenue streams and uncover innovative financial growth opportunities. Integrate data strategy with the business strategy, focusing on supercharging, streaming, and monetizing IP, with an emphasis on value cases over use cases.

- **Organizational alignment:** Take a commercial and value lens to the CFO, converting numbers into aligned objectives in annual revenue plans and targets for organizational alignment on data use for business growth. Move costs to business units to activate ownership of investment and value realization. Overcome organizational tension by aligning costs with where value is being realized, fostering shared responsibility and commitment across the organization.
- **Decentralize data skills:** Embed data scientists, engineers, and analysts into business units to foster ownership and trust, ensuring responsibility for delivering value lies with the business unit, not just the data team. Complement this decentralized approach with the center acting as the glue between different domains and organizational components. Adopt the “data as a product” route, introduce best standardizing practices, and implement a data mesh architecture to further decentralize control into the business unit.
- **Trio strategy development:** Establish a trio strategy involving the CDO, CTO, and business leaders, each with distinct objectives: business leaders for “last mile” adoption and value realization, a data unit for introducing a new narrative, and technology teams for reducing time to business value.
- **Guide technical teams:** Help core engineering teams understand the concept of higher-order assembling versus building, adopting a modular approach akin to a classic LEGO strategy for flexibility and adaptability. Establish a design authority community that conducts extensive horizon scanning through the Digital Advisory Council to preserve modularity in engineering and architecting. Initially rely on third parties to build teams with diverse skills like data engineering and data scientists, and develop a team model with a data product leader from a consulting background, who is business-oriented in terms of how they think and operate.
- **Bring the business along:** Focus on the value case and determine what needs reporting. Tell a story to help functional leaders reimagine the “after” and quantify that for the board and CFO to convey the commercial gain. Emphasize change management, solve the “last mile” first, empathize with stakeholders, and manage expectations by clearly communicating priorities. Develop an approach from experience and adapt existing methodologies to

the unique situation of each organization, understanding that there's no single playbook for this journey.

- **Future possibilities:** Recognize the implications of Web3 for data and digital-driven businesses, acknowledging the potential for cross-industry collaboration and innovation. Avoid common pitfalls such as overestimating adoption and underestimating the change challenge, and mitigate these issues through early embedding of individuals and clarity about the consequences of prioritizing specific areas. Showcase how other organizations with similar functional processes are managing these changes to make initial conversations more accessible and provide valuable insights for future strategies.



13

Nathalie Gaveau

AI Tech Entrepreneur and Board Member

I think AI will be much faster in some sectors like content creation or coding. So, the key is the strategic promotion of AI in business environments, educating businesses about its impact and encouraging them to pilot it for specific use cases... I have never seen any program that works straight away everywhere, so you need a phased approach, designing pilots and achieving small wins to get buy-in from the team. Probably everybody in this book is going to say this – but it's true. When you can say, “We increased sales by 15% in this region in a specific geography because we used X for the automated pricing and push” – people love that. Real-world success stories can help convince stakeholders and accelerate the adoption process.

LinkedIn: [nathalie-gaveau-ai1](#)

NATHALIE GAVEAU

Nathalie Gaveau is a serial entrepreneur, board member and investor in technology and innovation. She started at Lazard and then founded and scaled successful startups, invested in emerging companies, and served on the boards of leading organizations worldwide. The first eCommerce company she co-founded, PriceMinister, was acquired by Japan's Rakuten for over \$200 million in 2010. She served as a foreign trade advisor of France in the UK and was made a Knight of the French National Order of Merit in 2015.

LISA: What are the first things you do when you go into a business?

NATHALIE: I'm very curious so I start by asking a series of questions and learning. "What is the foundation for these analytics? What is the state of the customer segmentation? How's the distribution? How do I optimize my trade promotions and assortment?" Then I go into category management. "What are my consumer trends? What's the marketing return on investment and how do I position myself within my category?" Then I move to the supply chain, into things like dimensional inventory optimization and smart maintenance (if there are factories, maybe use a spend analytics product like Life Cycle). Or I move into product management and technology if this is a software business. I need to understand how coordinated the different teams are, and how customer-centric they are.

Personally, I like to start with the customer, focusing on customer segmentation and market strategies to identify opportunities for data-driven decision-making. Then I go to the market and the salespeople because they quickly grasp that if they have better data and can target the customer knowing how to make them more successful, then they will also be successful subsequently, so that's usually where I get the fastest response.

I've seen different-sized companies swiftly grasp the importance of business data analytics (for example, I'm on the boards of Coca-Cola Europe Pacific Partners, one of the world's largest FMCG, and also PortAventura World in Barcelona, one of Europe's largest amusement parks, which is a private equity-backed business) and they typically respond by creating data teams. If there's

nothing in place, they'll hire talent and learn fast how to use data. It can be within the sales and marketing team or the IT or digital team, but when it starts to develop, it becomes its own data team.

LISA: How have you progressed with data in your career?

NATHALIE: Starting in mergers and acquisitions in finance, data was always integral to my work, whether it was projecting revenues or analyzing sector data. I then co-founded PriceMinister, a French e-commerce platform that became a large marketplace, where data drove our supply-demand matching and informed our CRM strategies. Partnering with ex-mail order catalogs that had huge databases but didn't know how to use them, we became really good at segmenting and automating recommendations from different buyers, giving visibility to sellers on the sales and the pricing of their products. Eventually, this information became part of the DNA of the company and a tool to pilot other categories.

Next, I built a social platform, selling via social media that essentially rewarded influencers or content providers based on how many sales or views they generated – a slightly different way to do advertising. It was a powerful data platform that analyzed content-to-commerce conversion using advanced tech like MongoDB and Elasticsearch (a very good stack) and machine learning for tasks like fraud detection and automatic product categorization or tag matching. We had 150 million products so needed to standardize and auto-categorize products that were coming from multiple sellers; for example, an online article in ELLE magazine would talk about white shirts, so out of your 150 million product catalog, we'd match the most relevant products based on content but also other data available to us. Firstly, we needed to identify the focus on white shirts and potentially specific style attributes; then, for greater relevance, take into account the individual profile of the visitor, their past purchases and preferences. We did this on both written and visual content using Natural Language Processing (NLP) and visual recognition. This was very useful to monetize further content with new formats and generate a ton of shopping intent and revenue. As a consequence, we sold the company to a large media

group that used the technology to devise new content and advertising strategies for their clients.

Now, my career involves investing in startups including of course the intersection of content, marketing and commerce, but also AI ventures and impact (a super carbon-accounting firm). And, on the side, I serve on the board of four large companies including Lightspeed Commerce, a leading software company in the POS and payment world. Data and AI are pivotal for them to keep product leadership, for example, to help automate menu generation and analyze product performance as well as in customer service or development. With customers around the world, they need to optimize standards and performance. They can do this using multiple tools, for example, enabling them to train people or get really skilled English speakers who can interact in real time with people worldwide in multiple languages. I think it is only the beginning!

LISA: How do you start the conversation with a company if they're not used to working with customer service or training in that way?

NATHALIE: I've worked on both the cost and resource sides, so I know the hard jobs to find, the people that are hard to retain, and that training takes a time commitment. So, I usually show the business practical examples of solutions implemented by more advanced companies, allowing them to identify their own issues. There can be quite a lot of resistance to putting new solutions in place, so it's key to find champions within the organization who can manage these trials and iterate on the approach based on the learnings. I help them do a pilot scheme that may or may not work. If it doesn't, the important thing is to stop (not let it carry on...) and find better solutions. I think the most important thing is to see everybody and catch up, asking, "Did you implement this daily work? If not, what are the learnings?" I'm very iterative because of my background in developing tech products, and I think that's important. I am a problem solver, and I like pushing until things perform.

LISA: Where do you see the most challenges? Is it when companies are over-ambitious about using their data?

NATHALIE: Many companies discover their data isn't well structured, or enough for their needs, or sufficiently cleaned up, or they just don't own it as first-party data. The underlying challenge often lies in the data foundation itself, with issues like inconsistent collection methods and a lack of common analytical language. You need to establish this common language and build expertise in the companies with the right skill sets for data analysis, so there are multiple challenges but it's usually the foundation that has to be reviewed.

LISA: How should a company go about building an AI team?

NATHALIE: The journey starts with defining the use case and with the end goal: what do you envision as your ultimate goal? This makes you go up the food chain and reveals the need for a diverse set of roles: business experts, domain experts, product owners, data owners, business analysts, and data or analytics leads.

In a very large business, there are also commercial and supply chain analytics leads, for example.

Then you have the technology platform team, data architects, technology architects, solution architects, and those that operate that, with cybersecurity teams and infrastructure engineers inside it.

Then there are the data scientists who are coming up with the data products and who work with business and analytics leads to make sure the tech platforms provide them with the right data.

The first priority is building the hub and populating it with a mixture of these profiles to form the backbone of a strong data foundation – that's critical.

LISA: Can you share an example of effective AI usage from your board experiences?

NATHALIE: An instance that comes to mind is a large distribution company I work with. Despite its size, the company is very agile due to its well-established customer-centric approach and structure. They leverage business analytics tools across their value chain and are now enhancing their recommendation systems and capturing data from channel contacts with machine learning. Their power is to really master omni-channel data and predict your next order.

It's a large workforce, but they have more and more digital workplace integration. They use data in their supply chain to fulfill their major sustainability commitments, tracking water and energy usage. This company has millions of data records, including a lot of information on product range by customer segment, previous order history, customer basket size, price, and sales growth targets – this is how to effectively drive growth.

LISA: What are your thoughts on the adoption of OpenAI and large language models?

NATHALIE: It's a massive opportunity and I am very excited to see the emergence of more and more teams building AI-native products on some amazingly powerful stack! For example, I'm an investor in a company called Giskard that provides an open-source solution for AI quality, with a testing framework for AI models to eliminate risks of biases, performance issues, and security. I think companies in general are quite wary and don't necessarily understand the compliance and ethics around AI unless they're tech companies. Traditional businesses are very interested and flocking to it but are far from adopting it. Like any emerging technology, adoption and full-scale integration of AI will take time, mostly because you need to understand where to plug tools like GPT-4 into your own ways of working. The process reminds me of the initial hype around nanotech. While everyone predicted it would be the next big thing, in reality, it took a decade for it to be fully understood and used effectively. I think AI will be much faster in some sectors like content creation or coding. So, the key here is strategic promotion of AI in business environments, educating

businesses about its impact and encouraging them to pilot it for specific use cases, and using good governance and safety around it.

LISA: What would be your top tips for a company on getting value from data?

NATHALIE: Basically, data utilization should be strategic and organized. It helps pilot decisions and then coordinate execution. It's because of these two major management dimensions that I love data. It's really important to see how it is delivered in teams who are using it. It should be managed by a dedicated team of professionals who understand its value and can ensure its proper use. It's common to see tools like Tableau being used to create a multitude of dashboards that create a proliferation of information that is ultimately not used. This results in a waste of resources and outdated information. So, you sometimes need to be quite drastic with really good governance to ensure only relevant and actionable information is produced and used.

Often, you'll get some resistance, for example, from front line staff, who consider leads and customer data as their proprietary information. They might say, "Why would I give the information about a lead?" That's why, in my view, it's always harder to do data transformation inside groups than to add external plugins like new e-commerce or supply chain tools. It's much easier to add something extra than it is to change how people work and redesign a workflow or access information. But overcoming this resistance and instilling a culture of data sharing is essential for a data-driven transformation because it's not enough to have access to vast amounts of data; it's what you do with it that counts – knowing how to leverage it for actionable insights.

LISA: What's been an important lesson you learned from a failure?

NATHALIE: We once had a pilot program to automate product pricing that didn't yield the expected positive impact. The key issue wasn't necessarily the data but our understanding of the market and the consumers wasn't clear enough. So, this "failure" made us realize we needed to review the whole go-

to-market strategy and start again. It also highlighted the iterative nature of AI and data-driven projects. Instead of abandoning a project after an initial failure, the focus should be on continuous improvement and refinement.

LISA: What is the most impactful approach toward becoming an AI-driven company?

NATHALIE: I think it starts with a vision from the top leadership with the CEO leading with a clear strategy to transform the company into a data-driven or AI-driven entity and then establish a timeframe, so people realize this is going to happen.

A proper foundation of business analytics is crucial and takes time to build, often a year. Overburdening the team with a rushed transition might make people feel overwhelmed and result in high turnover. So, start from a strategic perspective, giving the business a timeframe to get the data foundation right, and gradually introduce AI into different aspects of the business. With these new pricing tools, some of the products were not even available, so you need to frame it in a strategic way and communicate with the organization about the ongoing changes and their benefits.

I also like a bottom-up approach to let people find smarter ways to solve their challenges; they become more curious about the tools, experiment, and often collective intelligence works very well.

I have never seen any program that works straight away everywhere, so you need a phased approach, designing pilots and achieving small wins to get buy-in from the team. Probably everybody in this book is going to say this – but it's true. When you can say, "We increased sales by 15% in this region in a specific geography because we used X for the automated pricing and push" – people love that. Real-world success stories can help convince stakeholders and accelerate the adoption process.

Key Takeaways

- **Extracting value from data:** Deriving value from data requires a strategic and organized approach, managed by a team who understands its value and can ensure its proper use. Merely having vast amounts of data is insufficient; leveraging it for actionable insights is crucial.
- **Early diagnostic for AI opportunities:** To assess how well data informs current decision-making across the organization, start by asking questions about key areas like customer segmentation, promotions, and supply chain. This initial exploration will quickly uncover high value opportunities for further investigation.
- **Phased approach:** Designing pilots and achieving small wins can help get buy-in from the team. Real-world success stories can help convince stakeholders and accelerate the adoption process.
- **Learning from failure:** Failures can provide valuable lessons. For instance, a failed pilot program to automate product pricing can highlight the need to review the whole go-to-market strategy and start again.
- **Dual approach to empowering the whole organization:** The successful application of AI requires clear executive directives in a top-down approach and a comprehensive change management program. Active participation from every level of the organization in a bottom-up approach is also essential. It's crucial to ensure teams understand the benefits of new tools or solutions, how to use them, and the improvements they can bring to their work.



14

Phil Harvey

Applied AI Architect

When you look at what's next in terms of synthetic generation, synthetic lead-generated data simulation, causal ML, and graphical methods, there's a raft of new tooling that is more about fitting together models of knowledge than it is specifically statistical models of data. I think we're moving from the data age into a new form of knowledge age.

LinkedIn: philipdavidharvey

PHIL HARVEY

Phil Harvey is an Applied AI Architect and Leader with 23 years of experience in AI. He is embarking on the next stage of his career and prior to this he was the Principal Program Manager at Microsoft. He is the co-author of the book *Data: A Guide to Humans*.

LISA: How did you get into AI?

PHIL: I was interested in computers from a young age, so I traveled long distances every day to study computer science because my local school didn't offer it. Then, despite my career counselor saying it was a very bad idea, I did a degree in artificial intelligence at Sussex University, which means I can say I've had 23 years of experience with artificial intelligence.

For me, artificial intelligence came first and computer science second. The universal Turing machine was designed to test out ideas of mechanical intelligence, then computer science picked up the computing side of it. Yes, I'm aware of the debt we all owe to Charles Babbage, an early pioneer in computing, and all that kind of stuff, but that's my story and I'm sticking to it!

After university, deciding I didn't want to be a programmer, I went into computer-aided design, helping an architectural visualization company manage their render farm to make artists more productive, enabling them to do 60 images a week as opposed to six. We didn't really think about it as data as people do now, but they had a 3D model of the entirety of London within the M25 as a managed asset and you could go backward and forward in time depending on the buildings that had been built. It was a surveyor's dream – a digital twin of London that managed to produce these images.

With the crash in the building industry in 2008, I moved into digital media – planning, buying, and advertising with an agency that had a file server in a mail server. As the sole technologist, I built machines, designed databases and data systems, and worked all that stuff out.

PHIL HARVEY

Three years later, I launched a data technology startup with the owners of the media agency, called Data Shaker, based on what I'd built – patented technology for mixing together time series for analytics. This coincided with the explosion in big data, where large datasets on human behavior could be analyzed by data scientists for the first time to reveal patterns, and the company thrived, but after five years I decided I wanted to give up writing code and take up talking to people and helping them instead.

So, I moved to Microsoft where I became a cloud solution architect for data and AI, spending four years talking to Microsoft Partners, which was very interesting. I then spent two years working in research and engineering the incubation of a simulation-based AI technology called Project Bonsai. It used deep reinforcement learning training AI models on synthetic data for industrial applications. At the time of this interview, I am just about to start a new role as a principal architect in the finance industry.

LISA: What kinds of business problems were you trying to solve with data in your different roles?

PHIL: The architecture visualization firm is the crispest example of a business problem. Their artists were very good at creating aesthetically pleasing images of buildings that didn't yet exist, but it came with a lot of admin – managing rendering, gathering data, and having 3D models. Artists are not great at admin, but they are great at painting, so I created software that reduced the amount of admin and increased the painting. My automation software is still running today with multiple millions of jobs done.

I regard my work at the advertising agency as the spark for all the other forms of data work I did. Data gives you new ways of knowing and new things to know. To work that out you have to ask the main question – “What can I know from this dataset?” Lots of data scientists will say, “Be specific in the question you're asking and then we can find the data to support an answer to that question,” but I don't really adhere to that. I believe the general question

is really important. Spend time working out what you can know from the data you've got. When you know that, you can start making it useful.

You can ask, "How do you know that automatically and regularly? How do you get the right people to know the right things to inform better decision-making?" But it starts off with that single question, not a whole range. If you want to find out what data can do for you, you have to be able to mix it together and compare different things.

This fundamental principle continued to guide me when I moved to Microsoft. There was a whole bunch of tooling and a whole bunch of data so the questions were, "What can we do with these? What can we make happen?" And, if building an application in Kubernetes or a new website, "What does this do for our customers?"

So, when looking at data and machine learning, the questions are, what can we know, when can we know it and how is that valuable? It's more about the way that you learn and how you understand the world than it is about how you provide functionality. The best examples of solved business problems are toward the end of the journey where we've made it a knowledge business not just a functional business. So, we don't just build our customers a website; we tell them things about how their business is doing. That's a very different kind of interaction and a different kind of business problem that you help people solve with data.

LISA: How crucial was analytics in the corporate strategy of the companies you worked with?

PHIL: In the startup we were selling to customers who had analytics as part of their corporate strategy so were actively looking for tooling to draw out insights and fuel data-driven decision-making. With Microsoft Partners it was the same thing, trying to provide a service solution to customers who want to do analytics. So, a lot of the time the challenge was in the applied AI space, aligning the current state of the client's data to their desire to leverage AI, which often required us to chart the journey in between. For example, they might need to

PHIL HARVEY

do some analytics before machine learning, as opposed to AI. Or they may need to migrate their data somewhere else. But that's why technical architects in data should be the most valuable people in it because, rather than writing code against the specification, they're tasked with evaluating the current state against the desired, and then writing and drawing up the building blocks to help you get there.

LISA: Could you always trust the data you were given and, if not, how did you tackle that?

PHIL: I have strong opinions about data quality. A couple of years ago I saw a survey that said data quality is now a C-level concern. No, it's not. This is an opinion born from employing expensive data scientists who have said, "I can't do my job because you've got poor data quality." What that actually means is you've asked them to do the wrong work. If you've got data quality to answer the questions, it means all the work has been done before. But if you're trying to work out if you can answer a new set of questions, then, of course, the data doesn't fit because you've not done it before – it doesn't mean the data is of poor quality, it just means that you have work to do. You need to gather the data or structure it in a new way. Data quality is more about knowing the current state of data and working toward enhancing it to answer new questions. So, instead of a CEO saying, "Oh no, we've got data quality problems, the business is in terrible trouble," they should be saying, "Why are the data scientists complaining about our data quality?" We need data architects and data engineers, people who are not purely trying to apply one set of tooling to problems that have been defined by somebody else.

LISA: What were the main hurdles you faced when working with data?

PHIL: Surprisingly, the major hurdles were not related to processing large amounts of data or the lack of technology. That's no longer a problem in the space.

PHIL HARVEY

The real challenge is a human not technical one, getting people to align on the questions we want to ask of the data and what we seek to learn. For example, let's say I've got a large number of transcripts from my call center, and I have a data scientist. We now have to frame a question that has business value that the data scientist can answer. Data scientists can answer with statistics and math using that dataset. All the tools, from language models to text processes to databases to object stores, all exist on-premises and in the cloud, so the technology is not the problem. The problem you have to solve is getting people together and saying, "These are the questions we want to test; these are the things that we want to know, these are the learnings we're going to seek." This entails respecting the learning process, including embracing failure as valid and valuable. So, instead of saying, "Oh no that didn't work," you can say, "Look what we've learned, for example, 10 things our data won't do" or "We've learned those tools don't work for our data and for these reasons we need to make these changes." Then people can come together and say, "Right, we're gonna change the process for the call handlers so that the data when it's recorded has these features in it. Then after X amount of time we will have the dataset that the data scientists can use to answer these specific questions." The biggest blockage is not adapting our approaches according to the data we have.

LISA: How do you facilitate cooperation between analytics teams and IT teams?

PHIL: I actually co-authored a book with Noelia Jiménez Martínez (*Data: A Guide to Humans*) on cognitive empathy, a skill I believe is essential in this context. It involves the rational act of using listening to understand the needs, feelings, and paradigms of the people you're working with. It can either be done by bridge people, such as architects, or it can be part of a data scientist's role. Because it's a human skill, it unlocks a whole amount of success where the technology just won't.

LISA: What types of organizational structures get the best out of people and data?

PHIL: Two approaches work well: diverse integrated teams or adept bridge people. What that means is you need people to work together. The first brings people together to experience each other's worlds and promotes understanding of different domains. For example, in the call center, if a data scientist has never been in a call center, no matter how smart they are they're specialists so won't get the reality, unless they experience it. When they do, they can understand how to ask the right questions in the right way to get the data they need.

The second approach can be just as successful, using bridging people to talk to the data scientist and, for example, spend an afternoon in the call center and translate back and forth. They are technical enough to talk the talk, but you wouldn't want them to write the code, while they can understand what the call center people are doing through observation, but they couldn't run it themselves. They are invaluable to a business because they can build trust on both sides and bring those pieces together. So, if you can integrate teams, do that, but if you can't, then look for those special people and pay them properly. It's not a junior role; they have to lead on both sides and that's hard and takes good experience.

LISA: Have you found it hard to make a business case for investing in data?

PHIL: I had the pleasure of lecturing on innovation at university a couple of times, not as an innovation specialist but as somebody asked to give their perspective. The business case for innovation is always an interesting one, because business case means money – the investment of time, effort, and resources to address a problem. The way you get that money, and you get the valid business case, is by telling a good story, creating a compelling narrative that justifies investment. So, you need to consider, are you promising them a 5% improvement in efficiency or are you promising to tell a CEO five things they didn't know about the business that will enable her to make better or new decisions? What does the CEO want to hear? And then who's investing in the solution?

PHIL HARVEY

Often the business case will fall into something highly measurable, but the measurement doesn't really mean anything as potential downstream problems must be considered. For example, you might say, "Buy our great tool because you'll get 10% more throughput on your factory line." But while it might be useful, you might be asking people to work harder which creates quality errors through tiredness. So, the business case is important, but you have to emphasize that it's going to be iterative and not all results are immediately quantifiable. The best place to start is to accept the fact that you don't know as much as you think you do.

LISA: Where have you found the most traction and utility in the application of data?

PHIL: The simple answer is always where processed data becomes something meaningful to people. Regardless of the automation, AI, or any tool, the endgame is to create an impact that is meaningful to people – whether it's in a call center, a factory, or a hospital.

For example, if your business is canned fish, you have to get fish and can them. If you're increasing the throughput of your factory, you need to ensure the boats are able to draw in the right level of catch, and so on. Each analogy will come back to what is the meaningful change for people who really understand the core business. It could be a red or a green light on something; it could be a left turn to change traffic flow through a site. It doesn't necessarily have to be a fancy dashboard.

LISA: What measures do you use to evaluate success?

PHIL: The biggest impact I have seen is always on what people learn and how it alters their way of doing business. Reduced costs are one particular kind of metric but often impact measures can be qualitative rather than just quantitative.

One of my biggest personal successes was when working with a digital agency that changed their business model from an aesthetic branding agency to one

PHIL HARVEY

where brand and technology meet. They were able to understand enough of the tooling to say, “We’re not going to be in-depth data scientists, but there’s enough out there that we can plug together for our customers.” They changed the way their business worked to grow new value. They had the confidence to do that because they knew what they had to learn, they focused, and their business grew, survived COVID, and did really well.

LISA: What advice would you have for people who want to start using data?

PHIL: Firstly, start by learning from the data you already have. You can use new tools, but you don’t always need new datasets. Regularly review what you’ve got and explore different ways of utilizing it. That’s always the thing that I regret from my time in startup land. I tried to do too much that was new all the time. You should experiment with what you have. If you’re only using your dataset for one thing, use it for five other things and see what happens. This needs an investment and innovation mindset as opposed to an optimized “business as usual” mindset.

Secondly, don’t believe the hype around new tools, remain skeptical, learn from mistakes, and don’t be afraid to innovate. So yes, focus on what’s important to you right now, but allow for experiments, because then you’ll be able to learn your way forward. If you feel like you’re standing still, challenge yourself to try something new, so believe the hype for a bit and have a go. But then stop believing the hype if it doesn’t work out. Don’t send good money after bad. If you learned something from it, the money was well spent.

PHIL HARVEY

LISA: Have you had a failure that equipped you for success later on?

PHIL: This may sound super granular, but I think it's an important point – if your dataset includes PDFs, all other problems will fade away as you try to solve PDFs because it's difficult and time-consuming to get data out of a PDF (and always harder than any tools promise it's going to be). So, any other way you can get the data that's not from a PDF is better.

LISA: What are you most excited about in how data can support companies?

PHIL: The exciting part is that while there's lots of data, the availability of tools and computational power means it's never been easier to have an idea and try it out. There are lots of things to learn; in fact, I often push Shakespeare plays into the new tools I've got just to see what happens.

I also think there are some super exciting things coming in the next phase of AI tooling, focusing more on modeling knowledge than just data. We're currently in the commercialization of the previous phase, using large language models but when you look at what's next in terms of synthetic generation, synthetic lead-generated data simulation, causal ML, and graphical methods, there's a raft of new tooling that is more about fitting together models of knowledge than it is specifically statistical models of data. I think we're moving from the data age into a new form of knowledge age.

It's not so much about what it is, as it enables people to do really interesting work again, because you're talking about modeling things that you know and testing them in more robust and rigorous ways. You're using data and machine learning and AI tooling to help you get there as opposed to thinking that some giant data-driven model will give all the answers at some point.

LISA: What would be your final message to somebody embarking on their data journey?

PHIL: Focus on people, people, people, people! Because tools in and of themselves are just that. But people make things happen.

Key Takeaways

- **Data as a new way of knowing:** Data is not just a collection of facts and figures; it provides new ways of knowing and understanding the world. Data can uncover hidden patterns and relationships to open up new areas of knowledge that were previously inaccessible and unknown.
- **Importance of asking general questions:** While specific questions can lead to specific answers, it's the general questions that often lead to the most profound insights. Respecting the learning process is crucial for gaining knowledge. This includes embracing failure as a valid and valuable part of the journey. It's through understanding what the data won't do or identifying which tools don't work for the data that we can make necessary adjustments in processes.
- **Solving business problems with data:** Deriving value from data involves transitioning businesses from being merely functional to being knowledge-based, using data to inform performance and enable a different kind of problem-solving. In both startups and established companies, analytics play a crucial role in corporate strategy where, frequently, the challenge is the alignment of the company's current data state with its desire to leverage AI.
- **Human challenges:** Often the most significant challenges in successfully implementing AI solutions stem from human rather than technical factors. Therefore, it's imperative to align people in the business on the questions they want to ask of the data and what they want to learn. Fostering cooperation between teams such as analytics and IT is vital and can be enhanced by cognitive empathy, which entails actively listening to understand the needs, feelings, and paradigms of the people in the business.

- **Importance of diversity in teams:** The success of AI projects hinges on the collaboration of multidisciplinary teams from various backgrounds with the expertise to extract meaningful data insights. Key to this synergy are skilled “bridge people” who can navigate and translate across different domains, facilitating cooperation and shared understanding between various teams.
- **Role of architects and engineers:** The roles of technical architects, data architects, and engineers are pivotal in bridging the gap between an organization’s current data state and its desired future state. Through a collaborative evaluation process, they assess existing data infrastructure and define the necessary building blocks to close this gap.
- **Data quality:** Data quality requires continuous improvement, not perfection. By understanding the current state of data, we can proactively enhance it to answer new questions and anticipate evolving informational needs. Data limitations are expected when asking new questions. Existing datasets may not perfectly align with previously unforeseen inquiries. This doesn’t reflect poor quality, but rather the need for data gathering or restructuring to address new needs.



15

Elizabeth Ajayi

**Director, Intelligent Industry at
Capgemini Invent**

Develop an ecosystem of partners. For large organizations, partnering with large technology companies can prove incredibly powerful. It's also essential to collaborate with specialists who have already been through this journey and understand what good looks like. Industry peers are an equally vital part of this ecosystem.

LinkedIn: [elizabeth-ajayi-robotics](#)

ELIZABETH AJAYI

Elizabeth Ajayi is a Director, Intelligent Industry at Capgemini Invent. She has led transformation programs for multiple clients across government and health industries, financial services and private sectors, and formerly worked for PwC and Accenture.

LISA: Tell us about your journey in AI and role today.

ELIZABETH: I began my journey in AI almost a decade ago, joining Accenture's AI and robotics practice. Being on the front line and seeing clients deliver value and productivity with these technologies showed me just how much potential AI and robotics had. After two years there, I moved to PwC in 2017 to help grow their AI and robotics practice until I joined Capgemini Invent in September 2023.

LISA: Describe your current role.

ELIZABETH: My role is client-focused and evangelistic. I assist clients with understanding AI data and emerging tech, and help them to devise a strategy and roadmap for adopting it responsibly, focusing on human-centered and transformative approaches.

Typically, my clients are very senior and C-suite level, so the initial support they need is understanding how AI is going to help deliver the strategic imperatives they have and how AI data and automation fit into the wider spectrum of digital tools and technologies available to them. We are working across a wide spectrum of industries and projects so we're seeing AI come to life in different ways.

While I was working at PwC, I played a key role in assisting public sector authorities in the thick of clearing the backlog inherited from COVID and also improving productivity to release more time for increasing demand. It wasn't just about tech delivery and implementation but about solving pressing problems, helping clients to understand and adopt AI so they can improve the outcomes for the public and citizens they serve.

LISA: What is the starting point in terms of identifying the challenges that you can solve with AI?

ELIZABETH: For local councils and public sector clients, the strategic priority is to deliver more with less, maximizing the benefits that taxpayer money can provide. This means identifying processes and activities that are consuming lots of time and resources. I look for opportunities to automate high-volume areas such as contact and demand management, such as the enormous number of calls that come into councils and the NHS about everything from house repairs and rubbish collection to a range of social care issues. By leveraging AI, we can reduce the effort and resources spent on routine tasks and enable staff to focus on more crucial issues, such as supporting vulnerable citizens.

At the same time, some clients aim to innovate, by launching new products or services to transform customer experiences. In these kinds of cases, we use AI to facilitate transformations, depending on the specific industry and problem at hand.

LISA: What are the main blockages you encounter in implementing AI to impact the business?

ELIZABETH: I've identified four main challenges.

First is the lack of understanding around AI, which often leads to resistance toward AI projects. To address this, you have to promote awareness, education, and showcase the potential of AI – the art of the possible. I try to show clients how I can help take away the mundane elements of their jobs and enable them to focus more on what they really want and need. With my healthcare clients, this is often about providing a human touch.

Second, there's the lack of mature cloud and data architecture. Many clients still operate on legacy systems that aren't AI-ready. Transitioning to the cloud is a massive prerequisite for AI and not having that can be a real barrier.

The third big challenge is that often there are outdated processes and operations that need to be streamlined before AI implementation.

Finally, there's what I call AI anxiety, where people think AI is going to replace them and so they don't want to engage in the program because they see it as a real threat. While AI might replace certain aspects of jobs, it's also set to create new roles and ways of working. Overcoming this attitude requires upskilling and reskilling people, enabling them to work closely with AI and focus on the human elements of their roles. If I can bring it to life, that really helps tackle this barrier. For instance, in the healthcare system, a huge chunk of clinician time is spent on administrative tasks that AI can take over, enabling the clinician to spend more time with the patient. By using AI in the right way to free up time and resources, you can literally save lives. You can be a lot more proactive and preventive with the interventions that you provide, and you can improve the quality of care people get.

LISA: Reflecting on your journey, what are the key lessons that would help someone just starting out?

ELIZABETH: Firstly, align your AI program with strategic priorities. It's important not to deploy AI just for its own sake. Instead, focus on how AI can help achieve your corporate strategy, mission, and purpose, and, crucially, make sure that any investments you make in AI are strategic and purpose-driven.

Develop an ecosystem of partners. For large organizations, partnering with tech companies such as Microsoft, Google, or AWS can prove incredibly powerful. It's also essential to collaborate with specialists who have already been through this journey and understand what good looks like. Peers are an equally vital part of this ecosystem.

Organizations need to pivot toward innovation with peers, finding ways to join forces to solve common problems and build on the progress made by others. The concept of partnership is critical to drive synergies and avoid siloed approaches.

And, lastly, engage from the outset, take your people on the journey, and make them the heart of it. People include employees, customers, and stakeholders; they are all fundamental to driving success. Empower them with the skills and knowledge about AI to be part of the conversation and deployment. Taking a human-centered and human-led approach is incredibly important.

Key Takeaways

- **Targeted AI investment:** Don't deploy AI just for its own sake. Ensure that any investments are strategic and purpose-driven. Identify time-consuming processes and look for opportunities to automate routine high-volume tasks such as contact and demand management.
- **Resistance to AI is multifaceted:** Several factors contribute to resistance toward AI projects. Lack of understanding and misconceptions of how AI can work, and its potential business impact, often leads to leadership apprehension. Legacy systems, not designed to accommodate AI, can create roadblocks, while outdated processes need to be streamlined before AI implementation.
- **Ecosystem of partners:** For large organizations, the development of a partner ecosystem can be a valuable strategic move. Collaboration with industry specialists who have traversed similar paths and have a clear understanding of successful outcomes is crucial. Organizations can also benefit from fostering innovation alongside their peers, seeking shared solutions to common problems, and building upon the advancements made by others in the field.
- **Adopt a human-centered approach:** Engaging employees, customers, and stakeholders on the purpose of AI integration early on is a key strategy for success. Equipping them with a roadmap, relevant knowledge, and the necessary tools fosters understanding and empowers them to participate meaningfully in the process. This collaborative approach paves the way for a smoother transition and, ultimately, a more successful AI implementation.



16

Louis DiCesari

**Head of Data, Analytics and AI at
Levi Strauss & Co.**

People trust things a lot more when they can see and touch and feel, but data science and AI don't always lend themselves to that. The trick is to empathize with your stakeholders and end users and figure out what's going to get them comfortable, confident, and excited, and then blend that with the analytical side of it.

LinkedIn: [louisdicesari](#)

LOUIS DICESARI

Louis DiCesari leads high-performing global analytics teams to deliver world-class data and AI products. He is currently the Global Head of Data, Analytics, and AI at Levi Strauss & Co, managing a team across the world. Prior to this, he was Group Head of Data and AI at Vodafone, and VP of Data and Analytics at Prudential Financial. He is a leading expert in building global data science teams across three very different sectors: fashion apparel retail, insurance, and telecommunications.

LISA: How do you start identifying the problem you want to solve?

LOUIS: In all three companies I've worked with, the need for data and analytics was recognized by the CEO and executive leadership. After the respective companies had a high-level roadmap, often from a consulting firm, I came in as an outside expert to expand and execute on the rough blueprint. Then, diving deeper, I used a 2 x 2 matrix approach to understand the value and speed to value.

LISA: What's the biggest challenge in your current role?

LOUIS: Working with physical product and supply chains is a new challenge. I find I'm happiest when I'm adding value in a balanced environment that juggles immediately monetizing projects with more significant, long-term bets. Some of our go-to-market cycles in retail are about 18 months long, so getting that balance right has been tricky in apparel, where I've had to ensure I have the right portfolio-based approach of things that will generate money in three to six months, versus bigger bets that take longer to materialize.

LOUIS DICESARI

LISA: What has your sponsorship from the company been and was analytics a focus of the corporate strategy?

LOUIS: At the beginning, I found that everyone was in support of analytics and AI but didn't necessarily have a shared understanding of what it meant. Education and diplomacy are always important steps to align the organization and get everyone aligned on what you're trying to do. One key sponsor was the Chief Commercial Officer. We recognized pricing and promotional strategy as areas that needed help as, until that point, it was done with only high-level insights, where things like price elasticity were static snapshots from years prior. This was coupled with the fact that we weren't making use of the very rich data we had both internally and externally. Another key sponsor is the Head of Merchandise Planning and Inventory Management, where AI can accelerate and streamline things like size curves. AI delivers smart starting insights so people can then address, "Where are the edge cases? The anomalies? The unpredictable pieces?"

LISA: How did you approach data quality?

LOUIS: I take a pragmatic approach. I focus on using the minimum data we need to start impacting a business problem and build from there. For instance, some of our external data on pricing didn't validate to a point where I was comfortable using it in models. So instead, we built models based on the internal data we had and then used planner and merchant judgment to calibrate on that. We could have spent six months trying to clean up the data and not have any results, or we could just move forward with something that uses smaller data but gets results faster – my bias is almost always toward the latter. The aim is to move forward with good-quality data, yielding quick results while continually improving the model.

LOUIS DICESARI

LISA: How do you collaborate with different divisions?

LOUIS: We have to remember that our job is to solve problems for the company. Be selfless with resources and funding and focus on the result.

LISA: Can you describe your team and its structure?

LOUIS: My team consists of data scientists, analysts, and data and machine learning engineers. We work in cross-functional squads focusing on specific domains, like pricing, demand planning, and eCommerce. This setup encourages domain as well as functional expertise. We loosely follow the agile methodology of project management which involves breaking the project into phases to ensure efficient incremental delivery, working in two-week sprints.

LISA: Have you ever created a product from the existing data, or a service that you then sold on to other customers?

LOUIS: Yes, we've made data products at Levi's such as price elasticity models and demand forecasting tools. The price elasticity models get put into a tool or an app that the planning directors, commercial directors, and so on can use. Similarly, our size forecasting is now in an app that feeds into the systems the planners use and pre-populates recommendations.

LISA: How would you describe the results and impact of the work? What were your measures of success and what can you do now that people couldn't do before?

LOUIS: Some aspects of our work are easier to measure than others. For instance, when I came to Levi's, I had an instinctive feeling we were discounting too heavily and too uniformly, and it was backed up by the data. So that was one of my first projects because it was easy to implement, and when you're new to a role, you want to show impact within six months. We did some quick A/B tests, which split testing to compare two versions, and it worked. Using Poland

as a test and holding the Czech Republic as a constant, we first approved the results there. Then, we scaled it to Germany, Austria, and Switzerland. When that worked, we expanded to all of Europe, and finally, scaled worldwide.

Designing a perfect A/B test was probably impossible or would have taken us forever. This method, with rapid testing, ensured results within a few months. I tend to work this way, with 80% reliability, repeat it on a slightly larger scale, and if you get a couple of positive results, just keep scaling incrementally.

LISA: Is there anything you would have done differently?

LOUIS: There are always things I would do differently! For one, I learned early on in a previous job that it's vital to involve finance early on and use an independent neutral group for your measurements. Recently, I adopted an objectives and key results (OKR) framework to clarify our goals and track progress transparently for everyone. Additionally, I learned to manage expectations better. In some cases, people expect AI to be better than human judgment, and in some cases that's not realistic. I wish I had gone out more strongly with "This is going to get us to 80% of the result very quickly" rather than the expectation of "It's algorithm versus human." Instead, we should be asking, "Where do we use an algorithm versus where do we use a human, and how do we balance the two?"

LISA: Have you had any experiences of failure when using data that equipped you for later success?

LOUIS: Certainly, I've had my share of lessons. In my insurance tenure, we developed a product fueled by consumer research and data that was a huge flop. Although consumers loved it, it was a miss with brokers – a critical misstep in a business where products are sold, not bought. You must ensure that the seller is as passionate about the product as the end consumer. Another takeaway was the need for directional accuracy over extreme specificity, especially with insurance agents. For instance, using a simplified traffic light system for AI recommendations instead of a specific rank-ordered list worked better. Sometimes

exposing too much information can backfire, so creating easily consumable recommendations that resonate with people's psychology is crucial. Simplicity trumps complexity in model implementation. There were instances in my career where I had to simplify a model for successful business implementation. The goal is to empathize with your stakeholders and end users, making data science and AI more tangible and trustworthy for them, and then blend that with the analytical side.

LISA: What are you most excited about next?

LOUIS: I'm always excited to make an impact through data in business results rather than focusing on fancy innovations. That said, the current landscape does offer a lot of exciting advancements. Generative AI is obviously the latest buzzword. Rather than the broad use cases, I think it holds promise in its narrow applications, such as product descriptions and customer service interactions. In fashion and apparel, computer vision, and particularly 3D, excites me. We've seen remarkable results using computer vision for things like item similarity. Plus, 3D rendering will enable customers to experience products much better online. This could dramatically cut down e-commerce return rates.

LISA: Any tips for a business that is embarking on using data?

LOUIS: Absolutely. My motto, borrowed from a consulting firm, is, "Think big, start small, scale fast." This strategy has always served me well. Accept the fact that nobody has any patience, and you need to demonstrate incremental progress to secure investments. It's about eating the elephant in small bites.

Key Takeaways

- **Leading AI transformation:** Empathize with stakeholders and end users, showing how data analytics and AI can benefit them in tangible ways, by solving problems or enhancing outcomes. Prioritize the value and speed to value of data projects, balancing short-term wins with long-term bets. Assess the impact and feasibility of each project, allocating resources and time accordingly. Work with physical products and supply chains where AI can accelerate and streamline pricing, promotion, forecasting, and inventory management. Be innovative, using AI to create new products and services that meet customer needs.
- **Data and AI best practices:** Data quality is not a prerequisite for initiating data projects, but a continuous process that evolves over time. Focus on using the minimum amount of quality data needed to start impacting a business problem and build from there, rather than spending too much time on cleaning and validating data. Collaborate with different divisions and ensure they have access to the data they require for their projects, and sometimes allocate funds within the company to facilitate this. Structure a cross-functional squad that focuses on specific domains, such as pricing, demand planning, and e-commerce. Working in two-week sprints, follow an agile methodology to ensure efficient incremental delivery.
- **Data products:** Create data products that can be used by other parts of the organization, such as price elasticity models, demand forecasting tools, and personalized customer experiences. Make them user-friendly and accessible, and integrate them into the existing systems and processes. Measure the results and impact of work, using independent neutral groups such as finance for validation.
- **An agile approach:** Adapt products and recommendations to suit the needs and preferences of sellers and customers. Focus on business results, rather than complex innovations. However, be aware of advancements, such as generative AI, computer vision, and 3D rendering, that can offer new possibilities and solutions for your industry or sector. Accept the fact that nobody has any patience, so demonstrate incremental progress to secure investments.



17

Vickey Rodrigues

CTO/CDO in Insuretech, Payments, and Healthtech

I've heard about initiatives that mean well but fail because the stakeholder is disinterested, or the output was not what they expected. It's important you take a capability approach, executing smaller projects before taking on bigger ones, to prove you can deliver the expected results.

LinkedIn: vickey-rodrigues

VICKEY RODRIGUES

Vickey Rodrigues has over 25 years of experience working with leading IT services, healthcare, telecom, and consumer electronic companies. His areas of expertise include AWS cloud, mobile payments, enterprise mobility, healthcare, and the blockchain. Vickey formerly worked for Samsung, Microsoft, Dell, and Vodafone.

LISA: What was your journey into AI?

VICKEY: Having worked on both vendor and end user sides, I've seen how technology has evolved into a key driver as well as a contributor for business value. This dual perspective has helped me understand both the hype cycle and value cycle for technology and data analytics as important tools to create new products or to bring new insights into the business.

LISA: Where do you start when trying to solve business problems with data?

VICKEY: When I was at a healthcare company, we had to leverage vast amounts of raw data about various co-morbidities and conditions for our patients. As it was during COVID, these huge datasets included blood work, medical records, data from the government, third-party vendors, and the lab itself. The challenge we met was bringing all this unstructured data together to generate health scores that helped us identify potential chronic health risks and trends within the data.

LISA: How did you get the sponsorship behind what you wanted to do and how did you make it a priority for the corporate strategy?

VICKEY: Our mission to improve health outcomes for our patients was sponsored by the Chief Medical Officer in the organization who ran a chain of primary clinics all over the country. Our focus was to drive the organization to think about outcomes as the real value of what we do for the community itself. As one of the main touchpoints for anybody with an ailment to reach out to

VICKEY RODRIGUES

doctors, we used data analytics to provide actionable intelligence that supported doctors and frontline staff in their clinical decision-making. This was crucial in creating disease management programs that targeted hypertension, heart disease, diabetes, and so on, where you could predict whether a patient would develop any of these conditions based on lifestyle, diet, or other factors.

LISA: Was this initiative led by a commercial company?

VICKEY: Yes, there was a very strong commercial angle. When we had line-of-sight on the data, we could have offered it to insurance companies to craft specific insurance products for this user base. Of course, post-patient consent and approvals were taken. We were able to offer a packaged insurance company product to add to our frontline stores, which had previously not been possible.

LISA: What was the quality of the data when you started working with it?

VICKEY: The organization had evolved from a manual-based system to a digital one, so we had about a decade's worth of data scattered in different formats, including machine data in PDFs and other formats, as well as unstructured files like X-rays and ECG reports. Our first challenge was to normalize and feed this into a data lake to clean it up before proceeding further.

LISA: That sounds challenging! Where did you start?

VICKEY: We started with a couple of minimum viable products (MVPs) with partners, and quickly realized that it required a lot of innovative thinking, which we didn't have in-house. So, we ran a mini hackathon where we invited a couple of partners to tackle the problem of extracting legacy data. This enabled us to identify a partner with the relevant technical resources and skills to help us.

VICKEY RODRIGUES

LISA: In terms of the process, what were the main blockages and challenges?

VICKEY: We had both business and technology challenges. The technical challenges were around data accuracy, and the ability to filter through a lot of data, which could be digitized or not. We then had to take that data to stakeholders (business users or external parties) to validate and confirm that these data models could be trusted and taken further to talk to customers about.

LISA: Did you need patient permission to use this data?

VICKEY: Yes, we needed patient consent when sharing with external parties, such as insurance companies. However, we masked all personal identifiable information (names, addresses, emails, IDs, and so on) and clustered it by age groups and location.

Our privacy policy as well as our data sharing policy allowed us to use data where it was targeted toward improving the outcomes of the patient. If you are my customer and I am using your data to predict a score of you developing any of the chronic conditions, then I could come up with an intervention plan. That was okay because we were helping you on that journey.

LISA: What advice would you give to people just starting to use data, and is there anything that you would have done differently?

VICKEY: Firstly, use cognitive services offered by various cloud service providers as this will help cut the time to get a solution to market by at least half. Secondly, securing strong corporate sponsorship is critical to success for anybody starting a project involving data analytics or AI within a business product or process. That's the norm for all the large-scale transformation IT projects now.

VICKEY RODRIGUES

LISA: Have you had any failures that equipped you for later success?

VICKEY: We had multiple stumbling blocks in the entire project, so I'd say be ready to handle failures. During my six and a half years, we worked with at least 10 different partners at any one time with different challenges across all of them. At least eight of them did not understand something clearly or couldn't see any business value returned but the two that did were pretty successful. Also, it's easy to rely on only one vendor or partner, but if that partner goes down, you're pretty much stuck, so don't rely on only one.

LISA: What are you most excited about being able to do with data now?

VICKEY: The significant increase in technology and computer capacity has driven innovation and made it easier and faster to roll out projects, for example, from setup to going live within 90 days. The computer capacity and algorithms available now with cloud service providers make it quick and easy to roll them and scale projects.

LISA: How have you built up your data analytics organization over time?

VICKEY: We have two sets of resources. The first was focused on data analytics for business reporting; i.e. KPI business performance dashboards needed to be provided to businesses as a standard practice. As part of the IT organization structure, I was focused on ensuring this team provided accurate information to the business. The second was focused on innovative projects, essentially mining through the information the organization is generating to create new data models and operate new products. This involves collaborating very closely with partners so you can operationalize these models at a rapid scale while minimizing impact on the day-to-day activities of the rest of the organization.

LISA: What is the one thing people should remember when working with data?

VICKEY: Equip yourself with technical capabilities and stay abreast of the market's evolving capabilities. There are multiple ways to skin a cat; you want the fastest and most efficient methodology.

LISA: How did you earn the trust of your stakeholders to invest in these projects?

VICKEY: It was a journey. I built my credibility gradually. I was able to talk to my medical officer and say, "OK, this is something that I want to get done and I need your support." I've heard about initiatives that mean well but fail because the stakeholder is disinterested, or the output was not what they expected. It's important you take a capability approach, executing smaller projects before taking on bigger ones, to prove you can deliver the expected results.

Key Takeaways

- **Value over hype:** Technology has evolved into a key driver and contributor for business value. Understanding both the hype cycle and value cycle for technology and data analytics is crucial for creating new products and insights. It's important to take a capability approach, executing smaller projects before taking on bigger ones, evidencing value impact.
- **Data-driven decision-making:** Using data analytics to provide actionable intelligence is crucial to creating targeted management programs. This approach enhances existing processes and informs the development of further improvement programs.
- **From raw data to targeted solutions:** Data availability opens up opportunities to develop targeted products or services tailored to a particular user base. The initial challenge, however, lies in addressing data quality issues. Scattered data in various formats needs to be normalized and cleaned up to ensure its usability for subsequent analysis and operations.

- **Scaling up with speed:** The increase in technology and computer capacity has made it easier and faster to roll out projects. The availability of algorithms with cloud service providers makes it quick and easy to scale projects. Running a mini hackathon can help identify partners with the relevant technical resources and skills to tackle the problems. Be ready to handle failures and don't rely on only one vendor or partner.
- **Leverage the cloud for faster solutions:** Cloud-based cognitive services can significantly accelerate your time to market for AI solutions.



18

Sean McDonald

**Former Global Chief Innovation Officer at
McCann Worldgroup**

Our concern was creativity and thinking about how we could leverage data and AI in terms of having a killer cultural insight or something that would allow us to do really amazing, award-winning creative work. We were also interested in thinking about creative ideas that had data and AI embedded into them or that became part of the actual solution. Clients benefited from us thinking in this way.

LinkedIn: [sean-macdonald-4467072](#)

SEAN MCDONALD

Sean MacDonald is a seasoned C-Level executive with experience in leading innovation and digital transformation. As the former Global Chief Innovation Officer at McCann Worldgroup, he oversaw the Futures Group, a team of specialists dedicated to driving the future evolution of the network and its clients' businesses. Prior to this, he was the Global Chief Digital Officer at the organization and has also worked at a senior level at Ogilvy & Mather.

LISA: Tell us about your role at McCann.

SEAN: I joined McCann in 2015. When I came into the organization, a lot of my work was about future-proofing the company. We're a professional services firm and we're constantly trying to stay at the cutting edge of emerging technology so that we can better service our clients.

Around 2018, I spent two years really focusing on data and AI because it was starting to get traction in the marketplace, and we thought we needed to take it very seriously. I was looking at how to upscale our organization and prepare it for this next wave of marketing that was inevitably coming. Right now, I feel like there's this new, subsequent wave with generative AI that's taken hold. My team has been focused on the Metaverse and generative AI and the emergence of Web3.

LISA: What types of problems were you trying to solve for your customers and how do you prioritize them and get them executed?

SEAN: It's interesting because they had their own internal challenges, which were primarily around having their data siloed across different parts of the organization. They were very protective of that data and didn't really know how to coordinate across silos or think holistically about how they could leverage their data. Their primary concerns were about understanding customers and then understanding the impact on the business.

Internally, customer insight was super important to us. I was using our datasets and also trying to have access to clients' datasets, which varied in terms of how willing they were to share information with us or not.

Unfortunately, sometimes they weren't as transparent as we would have hoped, because we could better service them if they were, but I think we were concerned with those two sides of the spectrum of customer insights and then the business impact of our marketing efforts.

Our other concern was creativity and thinking about how we could leverage data and AI in terms of having a killer cultural insight or something that would allow us to do really amazing, award-winning creative work. We were also interested in thinking about creative ideas that had data and AI embedded into them or that became part of the actual solution. Clients benefited from us thinking in this way, even if that wasn't top of mind for them. They were pleasantly surprised when we came up with insights and new ways of thinking.

LISA: How did you overcome the challenge of working with businesses that had siloed data?

SEAN: Because I've been working at a very high level in the organization, I wasn't dealing directly with clients who had these issues. But anecdotally I'd hear about it from the teams. I don't have many personal war stories, but there was a media company that was just sitting on a wealth of data, and they felt that for a competitive advantage they didn't want to share it with the creative agency. My personal view was that this was a real disservice. I think there was an industry politic there, with people trying to maintain competitive advantage. We were definitely all cognizant that if you held the data or if you could create a system where we were managing the client's data, then that ensured a longer-term relationship with them because it would be very hard to disentangle us and them. So, the dream was always that we could create our own internal technology where clients could port their data into it, and we would crunch it and analyze it. But we struggled to get that up and running because clients all had such different systems.

There was tension between the ad agency, the media company, and the client. The client wanted the data, not even knowing what their data was because it was fragmented across their fiefdoms. Then we were in our competitive fiefdoms with the media agency. So, basically, you have all this amazing data sitting in pockets, but there's no executive champion demanding that it all gets pulled together, nor a desire to do that because everyone feels like the closer they are to their data, the more they hold their data like it's a valuable resource that they should retain.

With some clients, you often have this disconnect between sales and marketing. Sales often devalue marketing. Marketing tries to drive sales to them, but they'd say, "Oh, we're better than you. We don't need your leads," and so they are certainly not going to be super excited to share their data or conversion data with the marketing team. That's a big disconnect.

I always feel like customer care is just somehow in its own world, so that's probably been very underleveraged because there's a wealth of consumer insights sitting, for example, within call center transcripts. We did some interesting work with one client where we went into their call center data and scraped it all and we were able to see what would garner the highest conversion.

There's also a separation between IT and marketing. There's been this back and forth of where the power resides. Is it IT who is dictating everything, because they're actually determining what systems are going to exist and how they're going to capture data? Or is it marketing who wants to leverage that data? Sometimes the key marketers are less technologically agile, so they don't actually have the language or the ability to truly analyze or give recommendations around what the CTO is doing.

I feel like there are so many dynamics at play. Often politics and people's personal motivations seem to work against integrating the data. I think at the core, people are not incentivized to share their data with one another.

LISA: How did you build momentum behind your work at McCann?

SEAN: There are three segments of the population. There are these very forward-thinking innovators who love to geek out on what's next and are always going to be the people calling me and attending the meetings and presentations where we're trying to introduce new thinking. Then there's a second group who really have ownership of the largest clients and the biggest pieces of the business. And, by necessity, they need to maintain a competitive advantage and retain the business, so they need the knowledge. They may not be as naturally excited about it, but they know that they need to get in front of it. So, they're very receptive, but maybe less enthusiastic. It's like bad-tasting medicine. And then there's a whole other set of people who are like, "I don't understand this. I'm not interested in this. I can survive without knowing." I think that's like 80% of the population or more. So, within our ecosystem, if we can make a difference with even just 15% of the group who are overseeing our biggest clients and make some inroads like just having a conversation around data and AI with clients and getting our data scientists and analysts in front of them. Having them think about what the success metrics would be from a quantitative perspective, not only a qualitative perspective, but those are also all wins for us.

I've been in the business of change for so long and have spent most of my career confronting resistance. So, for me it's about the micro moments, like if I get a sparkle in an eye or a little change, or one door opens, or I get one meeting with the client. We created this fun Data Monsters Workshop so we could work with clients to co-create ways for them to integrate data into various parts of their key processes. We got to do that with some clients, and I was just thrilled. And it turned into real ideas that went into market. But a lot of the battle was just getting to the point of being able to run one of those workshops, let alone manifest something from them. But we're in the knowledge business, so a lot of my wins are around getting people smarter on the topic, starting to adopt a new topic and then ideally transitioning that into actual work that goes to market.

We had success with a bank in the Middle East and a big retailer in the UK with our workshop. Sometimes it was useful because it gave our teams a clearer understanding of where they were going to face resistance with the client and how to have these conversations early on with key stakeholders. A lot of the wins could be very soft, like how do we bridge this conversation with this client?

LISA: Did you have a sponsor for this work, or did you have to drive it from the bottom up and create that sponsorship?

SEAN: I was the executive sponsor for all of this. Our marketing business is all about data, but the challenge we sometimes have is that because they're deep specialists, when they start speaking, people can get a little intimidated. A lot of what we've worked on over the last five or six years is making MRM (McCann Relationship Marketing) much more approachable for people. The co-CEO of MRM used to call me the data whisperer because she was like, "I don't know how you do it, Sean, but when you speak to people, they seem to get what we do!"

LISA: What were the challenges that you faced?

SEAN: A large part of me wants to say, "What barrier didn't we face?!" But the biggest barrier really has been people's time, attention, and openness and whether I can reach someone. Here's a funny little example – the global business leaders used to have a weekly meeting and sometimes we would get on their agenda, and we presented this Data Monsters Workshop as an opportunity for them. I remember in this one call I said, "OK, well, I'm looking for three volunteers. Who would be willing to try this out? Who's interested?" There's always one person who's immediately like, "Yes, I wanna do this." And then it's like tumbleweed. I used to have a joke with myself where I would say, "I'm just going to sit here until I get three volunteers. And I really don't care how awkward the silence gets." I would just sit there and sit there and then finally someone would say, "We'll do it." So that was really the barrier I was

SEAN MCDONALD

trying to overcome. Can I get your time and attention on a topic that you're uncomfortable with?

LISA: How did you build up your team?

SEAN: Our team was intentionally small. I functioned as the central point, securing resources within our agencies and connecting them directly with clients. These agencies were client-facing, and my primary objective was to generate enough business for them to expand their teams organically. While I may not have achieved the volume of client acquisition I initially envisioned, I successfully fostered client demand for the agencies' services. This structure allowed the agencies to scale their teams based on specific client needs.

LISA: Beyond the conversations and the workshops, did you develop playbooks or solutions that could be scaled in McCann?

SEAN: We did. We created this thing called Integrated Intelligence, which defined the offering to be presented to clients. This allowed us to see and understand all the technologies we had across the global network to enable consumer insights, tracking of business results, looking at programmatic media, and even looking at things like production optimization. Originally, we wanted to build a proprietary tech stack and we were partnering very closely with IBM. We had a proposal from them, but the leadership here really resisted the price tag to build it, and then when we went back to check the concept with some of the people who we thought were our biggest advocates and would actually be interested in this, and they said, "We don't think a proprietary solution will work because clients have such unique technology. We'd be better as consultants on how they optimize their technologies."

SEAN MCDONALD

Because of that, we ended up turning it into more of a methodology or a philosophy around how we saw data and created a set of products and solutions in terms of consulting services that could live underneath it.

Then we made sure internally that people understood the differences between what data sources are, what data manipulation is, what sources we already have licenses for that we could leverage, and then what tool sets we have in terms of data manipulation that we can leverage.

Then, within the media offering, there are a ton of customer insight tools that we had and proprietary research that we did. That was a great data source that clients couldn't get anywhere else in the world. So, we were able to create quite a nice story. It would have felt cooler if it had some underlying technology, but we also didn't want to make a big investment that we weren't ultimately going to be able to sell to clients. We modeled it all out. We had a business plan, and we had a pathway to get there if we thought we would have receptive clients. But that was the big question mark.

LISA: What insights do you have for people taking on a CDO type role within an organization today?

SEAN: For me, the most important thing is a shared language, because I think that if people aren't actually talking about the same thing, it's very hard to have a conversation. A lot of what we try to do is to create a very simplified way of looking at what an organization would need as best practice. In today's world, you must get into machine learning and generative AI and explain how those live as a powerful brain on top of all of the data that you accumulate. To me, it's mostly about change management, and co-creation is always the best means to get through resistance. I feel like if you can first get a shared language and then get all the stakeholders in a room and understand what their objectives are, then you can have them figure out how they think they could collectively use the data to achieve those objectives. From there you can start to prototype new ways of working or find small wins within that to start to transform their data ecosystem or ways of working. That's the best means of starting to create change.

LISA: Is there anything that you would have done differently?

SEAN: A little more hand holding through the process. Because we were just starting to perfect a lot of our change practices, we went through these distinct phases of, say, a focus on design thinking, and then AI and then data. Then we got really into change management, and lately, we've been into Web3 and future foresight, and I feel like a lot of our change practices were starting to develop right on the heels of a lot of the data work that we were doing.

I almost wish it was flipped because I think if I had more of those change management skills, then we could have had stronger conversion. We had the knowledge piece down, but we didn't really have the conversion side fine-tuned. That's why I think we could have done a lot more hand-holding and been a little more tenacious in terms of ensuring what happened. We followed up with people, but maybe that's where I feel like we could have done more. I think it was a little too easy for everyone to let it fall down the priority list.

LISA: What are you most excited about in regard to AI now?

SEAN: Mostly I'm scared about what's happening now. But I'm usually an optimist, so I'm kind of challenging myself to work out, "Am I scared or am I excited?"

Spinning it toward excitement, I think that we're in this really amazing next sequence of wondering how humans and machines are going to work as creative entities together. I think that the possibilities of generative AI are super fascinating and interesting, and I think it will create a lot of efficiencies. I hope that in those efficiencies it will also open up more human creativity. Just as creative people have had tools for years to inspire and bring their ideas to life, I hope that generative AI becomes a tool like that. Ultimately, I think you're going to need the human touch, some soul, to make creativity really resonate with people. I think it's pretty apparent when work has only been created by a machine. I've listened to a lot of generated music, and it just doesn't sound as good to me. I can tell it doesn't have humanity in it. Eventually, maybe this

is all going to get so good that you don't need humans, but I hope that it'll be more of a collaborative setup than machines taking over.

Key Takeaways

- **Apply AI to customer care communications:** The wealth of insights available in customer care emails, transcripts and feedback is often overlooked. Using AI to interrogate this data can reveal valuable insights for sales and marketing teams, such as what drives the highest conversion or customer churn.
- **Upskill marketers with technological abilities:** There's often a disconnect between IT, which determines the data capture and AI systems, and marketing, which needs to leverage the data. Upskilling the technological ability of marketers will develop their understanding of the recommendations from the IT function and how to use these new tools for hyper-personalization, content creation at scale, predictive targeting, and advanced automation.
- **Harness resistance and creativity:** To overcome challenges associated with individuals' time, attention, and willingness to engage, it is crucial to cultivate an environment that encourages open dialogue. Promoting creativity and innovative thinking leads to distinctive solutions and strategic advantages. Conducting co-creation workshops focused on customer insights and the impact of marketing strategies, promotes teamwork among stakeholders. Incentivizing data sharing across functions will reduce silo inefficiencies, realize the full value of the data and lead to better decision-making aligned with common business goals.
- **Engage key players:** To build momentum and drive business impact of AI adoption be strategic. Focus on the forward-thinking innovators and those overseeing the largest clients and business programs.
- **Focus on action and early change management:** Ensure AI outputs are translated into tangible business decisions and actions. Implement change practices early to accelerate the entire AI adoption process and business impact.



19

Julie Gray

Head of Data and Internal Systems at Agilio

One of the early difficulties we faced was setting the right scope for our data projects. We sometimes delivered things that were too complex or too sophisticated. We learned the importance of simplicity and clarity and adjusted as we went along. It was a lesson in restraint: not to get carried away with creating impressive outputs that overwhelm people. This taught us to focus on delivering clear information that supported actionable insights.

LinkedIn: [julieannegray](#)

JULIE GRAY

Julie Gray describes herself as “championing data to drive evolution” and has been instrumental in transforming how data is trusted and used in an organization. She is skilled in many disciplines within the IT sector, including sales, account management, business development, and CRM implementation. In business transformation her experience encompasses system and process implementation, data analytics capability development, and integration. She has worked for companies including Royal Mail, Servelec, and CSC. She is currently Head of Data and Internal Systems at Agilio.

LISA: Can you describe your journey to where you are today?

JULIE: My career began 30+ years ago with Royal Mail. I was part of an internal IT group of 2,000 people, with a role in service delivery management. After working there for 8 years, the IT function was outsourced to CSC, a global American-based systems integrator with circa 60,000 employees operating in 70 countries. After 14 years at CSC, across a number of sales and client-facing roles, I moved to Servelec, a HealthTech software company of around 500 employees operating in a single geography. My role was to implement a sales operation capability and adoption of the CRM. I then moved to a newly formed function of Enablement, as the Head of Data Enablement. Two years ago, I moved to take up a new position at Agilio Software.

LISA: Can you briefly describe your current role at Agilio Software?

JULIE: I lead the implementation of Salesforce (Sales and Service Cloud and integrated telephony) from requirements through to deployment. Now we have a single CRM solution that replaced three pre-existing systems. This was part of an overall business systems architecture deployment that also included finance and subscription solutions. We are now moving toward building a data analytics capability, using the data from the systems deployed.

JULIE GRAY

LISA: Going back to your work at Servelec, what were the biggest challenges you faced?

JULIE: Servelec was in the early stages of its journey with its private equity owner, seeking ways to maximize its business potential. With a value creation in place, we needed to have solid measurements against our targets, so we focused on understanding what drove our success. During my four years there, we transitioned from having very limited systems with poor adoption to implementing full-fledged systems with a unified process, a singular data language and structure, and a powerful business intelligence functionality. My team and I managed the data analytics function, creating dashboards that were accessible to our private equity owner so they could access our performance data directly. Everyone loved that it was the single source of truth, and it had the analytics.

LISA: What was the business problem you were trying to solve at Servelec and where did the momentum and push come from?

JULIE: Our starting point was a very ambitious growth target, so everything had to be about whether we could realistically achieve that. The momentum came from the clear business perspective of working with our executive and private equity teams on these targets. From a personal perspective, I focused on stakeholder management, working with the teams to ensure buy-in and collaboration. Instead of looking at the gaps, our approach was, “Let’s get things rolling, get some quick wins.”

We had a very structured approach, with a five-year roadmap, broken down into annual plans with visibility as our initial focus. We wanted to build momentum because if something took too long, we risked losing people’s interest.

We chose our tools, built a team, and started surfacing data, creating transparency and an actionable understanding of our current state. Although there were gaps, we remained focused on establishing a single, reliable data source without any peripheral fixes. This allowed us to start immediately and remain focused. However, I believe the key to success was not merely the data structure but

JULIE GRAY

having the right governance, attending the right meetings, and involving the right people. It was crucial to elevate the importance of our work beyond just a support function. We had to assert the importance of data visibility in our operations and, luckily, we had a management team who valued this and totally bought into the transformation. It was key to celebrate every little success and be positive in all communications.

LISA: How did you get everybody aligned?

JULIE: It was about knowing our audience and knowing what the business was trying to achieve – then using both these insights to demonstrate that data is not a separate function but an integral part of the business's success.

LISA: Were there parts of the business that you found harder to convince than others?

JULIE: Yes, the service area was initially a bit more challenging to convince, but the sales, growth, and commercial teams were on board very quickly. I leveraged my background in sales and business – I'd been where they were in previous roles – which helped to build trust with the teams. There was a strong sense of empathy and experience that brought understanding and credibility for the changes being implemented.

Although you need to sell into the senior stakeholders, the people you really have to get on side are the people in the operational, front-end roles, like the sales and service teams who have to input the data. We measured people's use of systems, even down to their adoption of the dashboards. Then when we built the story to say, "These teams had more of a success journey," it created a snowball effect because everyone wanted to be associated with the success of the project. I found it's worth pairing someone with data expertise with someone with business credibility to keep demonstrating that you are doing the right thing.

JULIE GRAY

LISA: What created traction?

JULIE: Our CEO was our number one user of the dashboards. He was literally using them daily, which demonstrated the value of the data. Our CEO's determination to make data-driven decisions set an example for the organization. In addition, we also provided insights and commentary to guide the conversation with leadership. For example, we shared data that demonstrated that we had accomplished 80% of the target, and also shared insights on the levers affecting this result, such as volume, speed, conversion of pipeline – the factors that determined the result.

LISA: How did you prioritize the work?

JULIE: We aligned everything to our strategic pillars. Every request had to be assessed in terms of its value to a specific strategic objective, asking people in the business making data requests, "What will you do with the information? Will that be something that you can take action on?" If a request didn't align with our key goals, we would challenge it.

LISA: Can you give an example of a metric that you aligned with a strategic pillar?

JULIE: Beyond sales metrics linked to growth, we also focused on measuring service, customer experience, and employee engagement for a holistic approach to evaluate performance management. We recognized the interconnectedness of metrics. We joined the dots between strategic pillars while not losing sight that all had a bearing on the sales target. This approach kept our strategy balanced and ensured support from the whole organization.

JULIE GRAY

LISA: When you started at Servelec, what state was the data in and how did you build credibility in it?

JULIE: When I started, our data was largely in a single source, but not a single source of truth. There were gaps and inconsistencies, but we were determined to make the data visible even though we knew it wasn't accurate from day one. The credibility was built by making technical fixes to the system to improve data quality and pushing back on the functions when necessary.

I remember an example of being challenged with, "Please don't take it personally, but the data in this report is wrong." I replied, "I don't take it personally, because it is your data." Having the confidence to push back was key because teams needed to accept accountability for data quality. We could put the governance in place, audit and create reports to visualize the data, but we needed our users to care about it, because if they care about getting the data right, they are likely to care about everything else in their role and about continuous improvement. We shone a light on the teams that had excellence in data management and tried to create a competitive spirit around getting data right.

LISA: How did that work on a practical level?

JULIE: Not unusually, we started with sales – everyone was interested in sales performance but once we reached a certain point, we decided to look at other areas. We didn't go to the nth degree to make sales reporting perfect in every aspect. Instead, we maintained a balance across all the strategic pillars. Not going too deep into any single area ensured we could keep the focus on the full picture.

LISA: When you first joined, did you have a team?

JULIE: I initially started working alone. My first major assignment was to evaluate and restructure our CRM, which was our main data source. After about a year, I hired two business intelligence (BI) analysts and we also worked with an external company. Despite the small size of the team, we delivered a

JULIE GRAY

significant volume of work, including running reports for all business functions, across sales, service, finance, HR, legal, development, compliance, and facilities.

LISA: How would you describe the impact of your work on the business?

JULIE: I believe the impact was wide reaching. Our approach was embraced across the organization; our reporting tools were used in every meeting and even employee performance reviews. We integrated a culture of continuous improvement into our operations, with data feeding into all governance and meetings. Additionally, we were transparent, sharing the data directly with our private equity owners. There was no data massaging; it was accessible to everyone who needed it, internally and externally. It was a really positive transformation and extremely well received.

LISA: What challenges did you face?

JULIE: Initially, it was daunting. Keeping up energy and optimism was essential, given that we started with subpar systems and data. However, the company's positive culture and the excitement for the project made the journey enjoyable and rewarding. Yes, it was demanding, and prioritizing tasks was key, but overall, the experience was so positive.

LISA: What key insights would you have for organizations at earlier stages of using their data?

JULIE: Context is key. Understanding the company culture and tailoring your approach accordingly is crucial. Align everything with the strategic goals to win support at every level. Embrace transparency and don't be afraid to go quick, even if the data isn't perfect. Acknowledge and credit all those involved in the journey. For me, it's about focusing on the soft skills and those stakeholder skills as much as the technical side of building data models. And remember,

JULIE GRAY

a beautiful visual can wow people, but it needs to actually be effective; don't get carried away with them.

LISA: Is there anything that you thought was a failure at the time, but actually equipped you for later success?

JULIE: One of the early difficulties we faced was setting the right scope for our data projects. We sometimes delivered things that were too complex or too sophisticated. We learned the importance of simplicity and clarity and adjusted as we went along. It was a lesson in restraint: not to get carried away with creating impressive outputs that overwhelm people. This taught us to focus on delivering clear information that supported actionable insights.

LISA: What are you most excited about?

JULIE: My university degree is in English Literature so, for me, the thrill comes from storytelling. Turning complex data into a clear, compelling narrative that prompts action is what I enjoy the most. A beautifully complex data project that leads nowhere is less fulfilling than a simple one that drives meaningful change.

LISA: What would be your tips for people trying to present actionable data?

JULIE: Start with energy and ensure everyone feels involved. Regularly sense-check your work by inviting individuals who aren't as closely tied to the data to provide feedback. So, if it's a finance dashboard, get somebody from HR or a different team to challenge you on the simplicity, asking, "Is it clear enough?" Finally, lead by example – you have to walk the talk. In our team, we held ourselves to the same standards we set for others, even employing user feedback and dashboard usage statistics to measure our performance. By tracking which pages of data were most visited, by which kind of role and so on, we could assess whether we'd hit the mark or whether we hadn't. And we continuously refined our approach based on that.

Key Takeaways

- **Strategic alignment:** Establish a unified process, a singular data language and structure, and robust business intelligence functionality. Align all aspects of your business to strategic pillars, pairing someone with data expertise with someone with business credibility to align with business objectives, enhancing the organization's decision-making and strategic development.
- **Phased transformation:** Take a phased approach to transformation, starting with establishing clear business objectives and robust systems. The first phase of the transformation establishes the right insights and reporting mechanisms to make data visible in operations. The second phase builds the momentum for change driven by a clear business perspective. The third phase scales the established systems and processes across various channels to ensure that the transformation benefits the entire organization.
- **Engage sales teams early:** To mitigate resistance and cultivate trust, quickly onboard sales and commercial teams as these typically handle crucial data input. Engage these teams to reinforce the priority of visibility and accuracy in growth performance data.
- **Data transparency and ownership:** Make data visible from the start, even if there are gaps and inconsistencies. Continuously improve data quality through process and technical fixes and establish a culture of data ownership. Foster a competitive spirit around getting data right to underscore the importance of not just having data, but ensuring it's credible and actionable, and can be used as a strategic tool for decision-making.
- **Holistic performance management:** Beyond sales metrics, include those on service, customer experience, and employee engagement for a more complete view of performance. Recognize the interconnectedness of these factors for a balance strategy toward achieving sales targets.
- **Positivity can fuel progress:** Maintain energy and optimism, especially when starting with subpar systems and data. Nurture a positive company culture with an enthusiasm for the project to make the journey enjoyable despite the challenges.



20

Peter Jackson

Chief Data and Technology Officer at Outra

You don't have to complete A before moving to B. You can work on multiple elements simultaneously, provided you're doing it correctly... In that way, other use cases can lift off at the same time or later on.

LinkedIn: [peter-jackson-aa518711b](#)

PETER JACKSON

Peter Jackson is the Chief Data and Technology Officer at Outra. Prior to this, he held a series of senior data roles at companies including Legal & General, and Carruthers and Jackson. He is the co-founder of the CDO Summer School which now has 900+ international alumni and the co-author of three best-selling books, including *The Chief Data Officer's Playbook* and *Data Driven Business Transformation*.

LISA: How did you start working in data and AI?

PETER: I started my career as a business analyst, where I honed my problem-solving skills, breaking down business problems and understanding the process flow and skills I still use every day. In 1998, as the internet took off, I became a self-taught full-stack back-end engineer. Then I became CEO of a company that developed integrated CMS and CRM systems long before the big players had thought about such a thing. We were sticking databases behind CMS systems to restrict logins for membership systems, and to direct personalized content for people in their web experience. However, I realized that rather than building the platform, the real value was in the data inside it, and how it could be leveraged to gain more customers or make predictions.

That's when I moved into data, working for several years for a boutique agency called Pagoda Mason. They were often integrated into the big four firms because they talked digital but not data, which was a challenge because these firms lacked a data strategy – nobody understood what that was then. So, it was all within digital strategy which meant that the CIO and CTO had their arms wrapped around everything.

My first opportunity as a data leader came from the Pensions Regulator, where the CEO, Lesley Titcomb, with a staff of 560, was embarking on automatic enrolment for workplace pensions in the UK. It's a risk-based regulator, and she had the flash of inspiration that the risk would be visible in the data, so the collection had to be absolutely tip-top, efficient, and automated. This data was coming daily from HMRC and the payee, and also from employers. Her words at the time, "We need to have the capability to find the signal in all the noise,"

PETER JACKSON

actually blew my mind – that somebody from a civil service background could understand these things. We did some really interesting work there and I also got sucked into doing stuff in the Cabinet Office and the Treasury, which has been a pleasure and a real honor.

Then Southern Water recruited me to address their data-related issues, as they were hounded by regulators who didn't trust their data. I was primarily tasked to increase customer satisfaction, reduce operation costs, improve data assurance and quality, and increase shareholder value. It was immensely helpful to have these clear targets and I would say it's one of the first things you need if you want to have any success – consider what you are being asked to do in terms of business outcomes. If asked to improve data quality, ask why? What is the business outcome that it is going to help?

Next, I joined Legal & General in a very ambiguous new role they created, Group Director of Data Science, with a remit to leverage and evangelize the power of data and data science across the whole group. Broad and global, it was a lot of fun with lots of opportunities but also frustrating because there were lots of challenges too. After 18 months, I took on a second role as the interim CDO (Chief Data Officer) of the investment management business which led to getting into a lot of asset management and delivering a £35 million data transformation for which I secured the budget and initiated.

During COVID, I worked for a year with Exasol, a competitor to Snowflake. Despite their struggles, it was an enlightening experience as I had used both of them and so knew the comparative benefits. It was fascinating to see how products were sold from a vendor perspective.

Later, I was approached by Charles Mackey, a well-known figure in the PropTech sector – he founded Hometrack, which is an Automated Valuation Model (AVM), and sold it to Zoopla in 2017 – to work with him on a new company. Now I'm Chief Data and Technology Officer at a company called Outra, which I bill as Britain's foremost data consumer behavior, data science company. We hold 5000 data points on 30.7 million households in Britain. All of it is GDPR compliant.

PETER JACKSON

So, I've worked at a big corporate level in Elgin, an investment management business. I've done utilities and experienced that very strange monopolistic environment. I've done public sector, and now I'm doing fast-scale startup – it's a pretty wide scope!

LISA: Can you think of an example of how you've used data and AI to fundamentally change the way the business runs?

PETER: The first thing I've realized is that, in the early phases, it's less about technology than anything else. You've probably got all the tools you need in an organization to make a good start and probably get a very, very long way. There isn't a magic answer to be had waving a wand and getting the latest piece of technology. If you've got Power BI that's fine. If you've got any form of database that's fine, just use it properly.

The second thing I realized is that it's actually more about data and AI literacy, culture, vision, and commitment. A lot of that has to sit with the exec team because when you're asking for budget, you need to have a sensible conversation with them to get their buy-in.

Then I think it's about having the right operating model, and this isn't technology, it's about having the right people in the right place interacting in the right way. And, if you're truly going to become data enabled, you've got to give those people the power to do it within the operating model.

An example of a data-driven initiative was at Legal & General where I created Data Science Bootcamp (it wasn't much to do with data science but I called it that to make it sound sexy so people would sign up!). We offered it to everybody. It was an 8-week course that gave employees a basic introduction to data, which eventually led to advanced analytics projects. I asked everyone to bring a use case which was fabulous because I got to see everything that everybody wanted to do. At first, it was very much about training them in how to identify a use case, how to create a problem statement, how to start collecting data. In other words, if you're going to move a dataset, you've got to understand where it's

PETER JACKSON

come from, who owns it, what you can and can't use it for, and what the quality of that data is for the outcome that you want to achieve.

The first cohort had 24 people, which I paid for from my group budget. By the time I left, we were running four concurrent courses. I designed the course with NTT DATA (an IT service management company) and they delivered it with some of my team and my resources supporting them. The program was so successful that business units were willing to pay for their employees to participate. For me, that was the measure of success. If they could convince their CFO or their line manager to pay £1500 – bingo! That is somebody saying that this is worthwhile.

LISA: How did you articulate that benefit back to the business?

PETER: We measured success through learning assessments, tracking projects' ROI, and presentations at our biannual data science conferences. The learning assessments we did right at the start meant collecting our metrics with data so we could prove some advancement or success. Then we did the learning needs assessment at the end, so each individual could assess whether they had improved or not. They had to relate their learning needs to their role and their business unit so they were identifying where it helped. We then did a longitudinal study over the following year, to prove the impact on the business. This was supported by the HR department. They loved it because they loved the fact that we'd set up those metrics and we'd done everything right from a learning point of view.

Tracking the projects that they took back into business looked at the ROI when they presented their projects to the wider business in our biannual data science conference. They were sponsored by Microsoft and Google and we went to their campuses together. Getting everybody off-site, including all the CEOs of all of the lines of business showed the community their support and how they valued what they'd done. I was really keen to get some of those products operational to actually prove the success of the bootcamp.

PETER JACKSON

Another example was at Elgin, where I went head-to-head with IT on some projects in the run-up to Brexit. There was a big project around fund liquidity that IT had put in a bid to the change management committee. I submitted my own data team bid – doing a low-code, no-code approach in half the time at a third of the budget. We got the gig and delivered it ahead of time, within budget. Proving the power of the possible in an existing use case was incredibly significant and really enhanced our credibility.

LISA: What do you think is the right operating model for a company starting to use data and AI?

PETER: It's definitely evolved. Ten, even five years ago, the CDO was reporting to the CTO or the CEO. It then evolved in many organizations as CDO sitting alongside the CTO, reporting to the CEO. I think you see that a lot, particularly in the big corporations. I'm getting late in my career now, so I can be objectionable and unreasonable and come out with radical ideas.

I believe the CDO should have a significant role in the organization, even overseeing technology-related decisions. For instance, when I first arrived at Outra, we had a CTO, a chief data scientist, and then me as kind of data and product leader. Now it's just me – the Chief Data and Technology Officer – and I run everything, from the network and issuing of laptops, to the maintenance of Outlook. This morning, I was listening to my data science team talk about how to use linear regression and twin neural networks in propensity models, and they were saying they were worried about the big migration we're undergoing at the moment because of maybe asking the data engineering team to do too much. I said, "Don't worry about that. You're all part of the same team. But sit next to each other and talk about it." This model enables us to really move at speed because I can make three calls, one to my IT manager to say, "Right, I need greater flexibility in our AWS resources. Don't worry about what you're seeing in the billing because I know what's going on." I can then talk directly to my Director of Engineering and say, "We need this. This is the plan for the next five days. This is what needs to go into the next sprint." And I can talk

PETER JACKSON

to my data scientists and say, “You need to talk to the rest of the team about what it is you’re trying to do.”

So that, for me, is the perfect operating model. But a lot of big corps are a long, long way from that, with data scientists sitting in an ivory tower, desperate to have their Power BI or their tableau reports built for them and not knowing where to go or whom to ask. You’ll have data governance singing a song that nobody is listening to, and you’ll have IT continuing to waste loads of money.

LISA: What advice would you give to a company starting to use data and AI?

PETER: I think the key point is that it’s not all about the technology but the operating model and how you’re running the business. Because we’re a small company, under 100 employees, we’re very flexible with our organization and methods. The structure we use, which I call a “layer cake,” combines different skills for each project. It’s led by the product owner who sits with the sales team, the marketing teams and the go-to-market teams, but also with data engineers, data scientists, all united in a sole mission – to deliver the product. That’s how you develop products and get them into production really, really quickly.

LISA: Who sets the priorities in your team?

PETER: As I sit on the executive committee, I work closely with the CEO, the COO, Chief Product Officer and Chief Commercial Officer to set priorities. I then translate these priorities directly into my own wide-breadth team. I make sure our priorities align completely with our Chief Product Officer, and work with our Chief Commercial Officer and CEO, to ensure we have the right practices and organizational structures.

PETER JACKSON

LISA: How do you anticipate and respond to change?

PETER: We haven't perfected our operating model yet, but the key is to be flexible enough to adjust. We're constantly looking at business process flows, ready to break them down and rebuild, if necessary, for instance, we've got a workshop coming up to address an issue with one of our process flows this week, of course, I'll say, "Well, do away with that, we'll automate all of it!"

LISA: How do you prioritize a use case?

PETER: In our book, *Data Driven Business Transformation*, Caroline Carruthers and I discuss a multi-dimensional data strategy. Many people wrongly believe that data strategy is linear, in other words, you start with data governance, then data management, management information, business intelligence, and finally data science, and that you can't have your data science until you've been through that journey. I fundamentally disagree with this. You don't have to complete A before moving to B. You can work on multiple elements simultaneously, provided you're doing it correctly. If you're going to do a piece of data science, do the fundamentals (data governance and management) right, and lay the foundations there for other projects. In that way, other use cases can lift off at the same time or later on.

LISA: Can you give examples of how prioritization works at Outra?

PETER: Prioritization is really driven by circumstance. At the moment, at Outra, we are completely commercially driven; so the prioritization largely depends on what is going to make us the most money, what's going to have the largest market impact, or what is fundamental to what we do. Sometimes, there are competing priorities, for instance, I had to convince my exec members about an issue I elevated as a priority around property listings from Zoopla and RightMove. I showed them that our models were built from data which covered only 12% of the market, which created a significant limitation. So, I set a target of above 90% of market coverage. We then needed to turn that raw

PETER JACKSON

data into gold dust, which we did by address matching. You can then suddenly build your models on 93% of market coverage with 78% addresses matched, compared to the previous limitations of 12% market coverage with 48% of those addresses matched. While this didn't directly earn us money, it improved the quality and scale of our data, enhancing our predictive models and having a knock-on effect on everything we do. We're going for hyper-optimization of features in models, so we needed to go back to the start and say "What's our raw material? What are we working with?" It required communication, education, and a convincing story. A lot of it's about storytelling. Trying to convince my colleagues that that was where to start, not to tinker around at the bottom.

LISA: What advice would you give to organizations struggling to balance data governance with other priorities and issues?

PETER: These organizations face a balancing act between needing to implement much-needed governance and addressing urgent problems, some of which aren't commercial, such as education and team diversity. My advice to them is to prioritize data governance while also tackling urgent matters in the right way so, as they move into other use cases, the foundations are already there. It's not for the data team to set the deliverables, that's for the business. We can set up an executive group to make the ultimate decisions.

LISA: How do you win over hearts and minds when you first go into an organization?

PETER: That's the killer question! When I was working for a sports association, we started with a data and AI (not digital) maturity assessment. It examined 12 points of the data ecosystem, only one of which is technology, the other 11 are things like behaviors, frameworks, policies, governance, leadership. I played that back to the deputy CEO and the CEO, and the common denominator that emerged across all the business units was that there was a lack of trust in the data due to poor data governance. So, I showed them how good governance could address this trust issue and prevent conflicts over figures, using specific case studies. In my previous role at Elgin, I got data issues onto the enterprise

risk register, and data governance was the solution. All of a sudden you get the Chief Risk Officer and the CEO very interested in how data governance can solve this risk and this problem.

Key Takeaways

- **Role of Chief Data Officer:** Consider the Chief Data Officer as a significant role in the organization, beyond just managing data. This role can be a strategic accelerator, leading the transformation of data into value to drive the organization's growth and innovation, managing risks, and enhancing agility and competitive advantage.
- **Building a data-driven culture:** Emphasize data literacy, foster a culture of data-driven decision-making, have a clear vision and commitment. These considerations are more important than just focusing on technology and need to be championed by the executive team, especially when budget allocations are being discussed. An effective operating model involves having the right people interacting in the right way. This includes having data scientists, IT managers, and directors of engineering working together as part of the same team.
- **Prioritize data governance:** Address poor data governance urgently because a lack of trust in the data can lead to internal friction. Demonstrate to the Chief Risk Officer and the CEO how data governance can be a solution to data risk. Prioritize data governance alongside addressing urgent commercial matters.
- **Internal training:** Empower employees and enhance change management with a program of training. Start with a basic introduction to data and lead to AI projects. Train employees on how to identify a use case, create a problem statement, and start collecting data.

PETER JACKSON

- **Setting priorities:** Collaborate closely with the CEO, the COO, CPO, and CCO to set priorities, which are translated directly into the team. Be flexible, continuously review business process flows, and redesign if necessary. Prioritize according to real-world circumstances, such as what's going to make the most money, what's going to have the largest market impact, and strengthening the core foundations of the business.
- **Power of storytelling:** Tell a compelling story about what is possible to inspire colleagues and prioritize issues.



21

Mark Beckwith

**Director of Data Governance and Architecture
at the Financial Times**

AI has brought us to an interesting point. As we leverage this new technology, we need to balance our ambitions with mitigating potential risk via well-thought-out governance. An incredibly valuable byproduct of the emergence of AI is the increased awareness of the importance of establishing robust and reliable foundational data capabilities. I'm looking forward to helping advance our organizational culture around data, part of which will be to challenge the perception that governance is just about compliance. In fact, it's core to value generation because it's intrinsically linked to the data value chain.

LinkedIn: [mark-beckwith-9603551](#)

MARK BECKWITH

Mark Beckwith is the Director of Data Governance and Architecture at the Financial Times, having joined the business in 2012. Prior to this he was the Head of Database Marketing at Domestic and General and has worked extensively in data and analytics for companies including Royal Mail and Ford Motor Company.

LISA: Tell us about your team at the Financial Times and how the business model has evolved.

MARK: The FT Group, part of Nikkei Inc., employs more than 2,300 people worldwide, including 700 journalists in 40 countries. It includes the Financial Times, FT Specialist, and a number of services and joint ventures. Most of our employees work in London, New York, Hong Kong, Manila, and, more recently, Sofia in Bulgaria, which has become a tech and data hub since being established in 2018.

Today, our revenues from content (via subscriptions) or advertising are weighted toward digital. We have a paid circulation of just under 1.2 million across digital and print, including those who pay for individual access and corporate licenses for organizations. As pioneers of the paid subscription model in the early 2000s, we took a metered approach, offering a free sample of limited content before having to pay. That model has, and continues to, evolve, and includes annual and monthly subscription packages, as well as low-cost entry trial subscriptions that move to full-price subscriptions after a short period.

LISA: How has the data function evolved?

MARK: When I joined the FT 12 years ago, data capabilities were dispersed across multiple teams, which included a data engineering arm, a business intelligence team that focused on reporting, web analytics, an emerging data science capability, as well as database marketing, which I joined to head up. Increased demand for data resulted in growth across all these disciplines then, when Nikkei acquired the FT from Pearson in 2015, the director of the data

area joined the board as Chief Data Officer (CDO) with the remit to develop not only our data capability but to also spearhead the business becoming more data-driven. One example of this was putting data and insight into the hands of journalists. This was a novelty then but now it's a very different story, with roles in the editorial team focused on using data to better understand our readership and how content is cutting through to particular cohorts and demographics, such as female readers, which continues to be a focus for editorial.

Back in 2015, our CEO John Ridding sought to get the whole business behind a north star metric. We created an engagement score, which combined recency and frequency of reader visits as well as the volume of content consumed with the premise that increasing engagement with our content likely translates to long(er)-term paying subscribers.

LISA: How did the engagement score influence the company's operations?

MARK: We set targets based on the number of engaged readers and these metrics became integral to our operations and a component of our company's bonus scheme. Getting this embedded across the whole organization took roughly three years of immense effort and dedication as we needed to establish the metric, communicate it across the business, and get it baked into people's objectives and targets. With the engagement score universally understood and supported, we then found ourselves regularly attempting to translate this insight into actual value. How much are these highly engaged users worth to us in pounds, dollars, euros, or other currencies? This is how LTV (Lifetime Value) was born – which is an estimation of how much revenue a customer will generate for the business over the life of that relationship. Also called customer lifetime value (CLV, or CLTV), this is a critical metric for a company trying to gauge the cost efficiency of acquiring new customers and supporting them over time. It built on our engagement model by including a monetary value dimension. For example, over a three-year period, how many of these people were likely to stick around and how much would they pay over that time, and *voilà!* – you have LTV. Any new product initiative was then assessed using this new LTV framework.

MARK BECKWITH

LISA: What are the key challenges that you face today?

MARK: My team's key challenge is to maximize the potential to create value from our data. By engaging the business, we aim to make continuous improvements to the quality, reliability, and useability of the data that we generate or procure. This is essential in order to meet the FT's current and future needs. The team's efforts are critical to the future success of our data and analytics capabilities and, therefore, the organization as a whole. It involves working closely with product and technology, commercial, editorial, and legal colleagues at all levels in the organization.

A practical example of this is developing a comprehensive understanding of the data we have within the FT – data about our data. Knowing what data we have and where it lives is great but how we extract value from this is by having a clear understanding of how that data flows through our organization and how it powers our products and processes. I was just writing a paper today about the benefits of data lineage in our product development process.

LISA: Can you expand more on the data lineage challenge?

MARK: Companies now have an insatiable appetite to get more and richer data, which is great, but it often happens in a very organic and unstructured way, which we have been partially guilty of. As a result, we've definitely got an opportunity to up the ante in terms of improving our data maturity, in this case, how we acquire and utilize our data across the entirety of our organization. For instance, we didn't have a data catalog, which makes life for anyone wishing to get access to, use, and understand our data, much harder!

In January 2023, when the news was full of the ransomware incidents suffered by The Guardian and Royal Mail, this increased an already heightened awareness of a potential future data breach. This turned out to be quite helpful in promoting the need to have better data about our data, held within a centralized data catalog. While we've made a good start, tackling risk, people will really sit up and take note when it starts to deliver value. This is where the importance of lineage comes in. Being able to trace how our data flows through our organization

MARK BECKWITH

will help us bring products to market quicker, whilst also providing a means to “impact trace,” in other words, to understand the impact of changes on our data ecosystem, so we make those changes with our eyes open in terms of the cost of those changes.

LISA: How have you been tackling this challenge in data lineage?

MARK: As I mentioned, we’ve now had some really good traction in establishing a centralized catalog of our data, and the explosion of AI has really helped here, as the old saying of “garbage in, garbage out,” has never been truer. So, I am leveraging this in terms of helping to make more people aware of the need to ensure that the data underpinning our AI efforts is reliable and of high quality. We’re not unique in the sense that the journey our data takes, from source to report, involves multiple systems and teams. However, if you think of the whole end-to-end data chain, we’re only as strong as our weakest link, so if you’ve got one team with stellar data governance practices, it’s futile if others don’t have anything. So, my team has been focused on developing a holistic capability to enable us to map the entire data journey. We’re now beginning to see significant progress with this first step.

LISA: How important was the sponsorship of what you were doing from a leadership perspective?

MARK: It was vital. The presence of a Data Governance Council, established in 2017 (largely made up of board members) has been a huge help. One initiative in particular that has benefitted from their sponsorship, has been the creation of an always on/BAU data mapping capability within my area. It was something that both my team, as well as compliance, felt was worthy of incremental investment, to ensure we had a systematic way to address any blind spots we had in terms of what data we had, who owned it, how it was being used, and so on. As we continue to develop our data governance capabilities, that sponsorship will remain vitally important, especially when trying to address any potential gaps in data ownership/stewardship.

MARK BECKWITH

LISA: What would be your key learning from working with data and AI?

MARK: One thing I learned quite early on at the FT from which I got a lot of confidence and traction, was seeking out advocates in unexpected places. For instance, our B2C marketing director, not closely tied to data, became instrumental in pushing our data-focused agenda. This happened when I provided a high-level, accessible piece of analysis of the impact of replacing one propensity model with another. Much of this analysis was caveated with some assumptions thrown in, but she said, “I think your best guess is better than anything else we’re gonna get, so let’s go with it.” I thought, “Wow,” because the effort that I’d put in prior to that in terms of making the work understandable and focused on what she cared most about, really paid off.

Another lesson I learned is the power of simplicity in explaining complex concepts to make them accessible, which is vital if you’re going to get people to come on the journey with you. So, I consistently emphasize to my team the value of repetition and simplification in communication, regardless of the audience. Some people may be uncomfortable with simplification because they fear it makes what they do sound easy. But, in fact, it’s quite the opposite.

There’s an Einstein quote that the definition of genius is taking the complex and making it simple. You might be really smart and know all the technical details of something, but can you describe it to someone who doesn’t really understand it? That’s my mantra.

LISA: What are you excited about next?

MARK: AI has brought us to an interesting point where, as I said before, it has shone a light on the importance of ensuring the data that feeds these models is up to scratch and well understood. We’ve had some good early traction on this with the development of our data catalog, but we’ve got a lot of runway ahead of us. This is exciting as it is a multi-year challenge that I feel will be transformational for the business and, from a personal standpoint, something that will enable me to continue learning and developing.

Another thing I'm excited about is tackling the misconceptions about governance and demystifying the notion that governance is all stick and no carrot that will slow the business down and cost more. I'm looking forward to helping advance our organizational culture around data, part of which will be to challenge the perception that governance is just about compliance. In fact, it's core to value generation because it's intrinsically linked to the data value chain.

Key Takeaways

- **Data and AI function evolution:** Support the data function's evolution from fragmented silos to an integrated hub-and-spoke model. Establish a lean center of excellence to guide and optimize data and AI practices, ensuring efficient resource allocation and strong data foundations. Empower data-driven decision-making by embedding hands-on analytics teams directly within business units. Incentivize continuous improvement by acknowledging and rewarding the data and AI function's achievements.
- **Data catalog and lineage:** Prioritize the development and maintenance of a centralized data catalog to enhance data quality, security, and compliance, and enable data lineage to track the origin, transformation, and usage of data.
- **Data governance council:** Establish a data governance council with the authority and resources to balance the potential commercial benefits with data privacy, security, and ethical considerations. Align data practices with your organization's values and reputation.
- **Data mapping:** Map data to mitigate risks such as breaches and losses, and lay groundwork for a more secure and efficient data environment. Identify and document all data sources within your organization, along with the flow and transformations throughout the systems. Support the data function in securing the necessary funding and resources to complete this task.

MARK BECKWITH

- **Invest in data engineering:** Invest in a strong data engineering team to build and maintain a data platform that handles collection, analysis, and reporting efficiently. Align and integrate this team with the data science and insights teams to ensure effective collaboration to deploy AI solutions.
- **Power of simplicity and repetition:** Communicate complex concepts effectively, using simplicity and repetition, and seek out advocates in the business.



22

Kshira Saagar

Data Science at Wolt and Doordash

Data science is 85% change management and 15% actual technology. It's becoming easier and easier – anybody can pretty much build an algorithm and there are all kinds of generative AI coming – but getting human beings to believe in it and use it is the hard part.

LinkedIn: [kshirasaagar](#)

KSHIRA SAAGAR

Kshira Saagar has worked in analytics and data science worldwide. He has built global teams and created data and AI solutions in retail, banking, consumer goods, telecom, and insurance verticals for leading Fortune 100 clients. He is currently DoorDash's first Head of International Data Science. Prior to this, he was Latitude's first Chief Data Officer and served as the inaugural Director of Data Science for THE ICONIC, part of the Global Fashion Group. Besides his corporate engagements, Kshira serves on several advisory boards for AI and data startups.

LISA: How do you see your role in data and AI?

KSHIRA: I used to call myself “the lost and found guy.” I try to find answers using data through algorithms, dashboards, analysis, and more, to help people make decisions. Obviously, people have been making decisions with less data and largely based on assumptions for years, but now, just making one small, smart decision based on accurate data is going to give you the edge and allow you to do smarter things. I like to think of myself as a builder of teams and trust because you can't work with data alone. You need smart people to create value.

LISA: What is your role now and how did you get here?

KSHIRA: I started as a consultant solving point problems using data science. Then I worked with various companies, including Gilt Groupe, Zappos, and Marks & Spencer in the UK, solving challenges in website functionality and email personalization in the 2000s and early 2010s. Then, when I moved to Australia, it was like being in a time machine with a five-year lag, so I found myself being asked to repeat everything I had done years before, so it was fun. Then I worked as the Chief Data Officer for a bank, and now I'm with DoorDash, taking care of their international markets outside the US, which is 27 countries.

LISA: What business problems were you trying to solve?

KSHIRA: Most problems focused on making customers' lives easier. For example, in Marks & Spencer's case, we aimed to enhance the online shopping experience by showing customers the right products or guiding them to the right page, or sending only the most relevant emails at the right time. We also worked on operational efficiency, in areas like supply chain and warehousing, asking, "How do we save money doing smarter things at scale?"

LISA: In the companies you've worked for, was analytics always a part of the corporate strategy or were you a sideline trying to get in?

KSHIRA: Early on in my career, I was lucky enough to work in places where people recognized the need to do things differently and take a data-driven approach, but the challenges often lie in using new technologies effectively. For instance, I worked in one company where they had spent tens of millions on a massive automation to run the warehouse but didn't know how to utilize it. Buying expensive technology is not fun. You could buy a Ferrari, but if you don't know how to drive it, it's going to sit in your garage. And that's pretty much what that automation did. They threw it at the data people to figure out, probably thinking it was an unwinnable challenge. But that's the kind of challenge I love. So, my approach was, let's bring in a set of smart people, and figure out how to use the automation to scale picking and packing of products. And we did it – 40 times better than a human could. I often find I get a problem that management thinks is a hopeless case. To prove this wrong, it really makes a difference when you have a seat at the C-suite table versus when you eventually get given the leftovers and asked to solve the problem.

LISA: When you've joined companies, have you found the data foundations already set up and what was the quality like?

KSHIRA: It's been a spectrum of everything across projects. Sometimes the data was muddy or unclear, or people didn't trust it. My approach was to peel back the layers and find the valuable nuggets within. I'm a big believer in "done is better than perfect" so I start with what we have and make progress instead of waiting for perfection. By demonstrating the potential on the back of an envelope, we often gained support and resources to improve further. I believe everyone wants to solve a problem; they sometimes just don't know how to go about it. If you tell them, "This is where you need to go and this is what I need to get you there," then people are more than willing to come and help you.

LISA: Can you tell me about your work process and what resistance you have encountered?

KSHIRA: Technology is not the challenge; trust is the challenge. So, in the beginning, it's more about convincing people that this is a problem that can be solved with data and getting their buy-in because the technical hurdles can be overcome through iteration. We've had success in creating effective algorithms, but the real challenge lies in convincing people to adopt their output. For instance, we designed an optimal pick path algorithm in a warehouse that would use the latest OR-Tools (Operations Research Tools), with a 25% saving on resources but, due to its lack of intuitiveness, it faced pushback from long-time workers. For example, the supervisor said, "I've been running warehouses for 25 years, don't tell me how to do my job." To win them over, our data scientists would walk with them, show them the model, and incorporate their inputs into it. The ultimate goal is for the stakeholders to view it as their algorithm and their model – a tool that they understand, trust, and can effectively use.

In essence, data science is 85% change management and 15% actual technology. It's becoming easier and easier – anybody can pretty much build an algorithm and there are all kinds of generative AI coming – but getting human beings to believe in it and use it is the hard part.

LISA: What's been the biggest challenge that you've faced?

KSHIRA: The more human interfaces involved, the harder the problem. An industry such as fashion, with its highly subjective and artistic character, poses the most challenges. Trying to tell the most creative artistic buyer how they should buy clothing is tough. We try to resolve this by augmenting rather than automating, emphasizing that we are there to assist them by providing data-backed recommendations. So, we'd say, "We know that with this particular shirt, this is how deep you should buy, this is when it will run out, this is how long the lead time is and therefore this is why you should buy it. We recommend you buy this, but what do you think?" It's a question of figuring out how to work with people who are not technical or data-driven and get them to come along with you on the journey.

I think the most important thing is understanding that the process is not merely about handling numbers or algorithms. It's about transforming these into meaningful actions that can drive a business forward. Leveraging data and AI in business operations demands a unique balance of technical know-how, creativity, and the ability to work with a diverse team.

LISA: In your past projects, did you start from a business case, or did you have to build it?

KSHIRA: Starting from a business case, instead of a data or technology problem, can set the project on the right path from the start. Data should serve to answer crucial business questions and facilitate decision-making. Our starting point was usually known inefficiencies that needed to be addressed. While this gave us the freedom to tackle the issues, it also presented the challenge of defining the precise nature of the problem. As we progressed, we realized that using tools was not necessarily a superior solution; the real metric for success could be achieving the same result with fewer resources.

LISA: How do you get the right resources behind you and build a team?

KSHIRA: My approach is akin to a “freemium model” – demonstrating the potential of a solution with a small-scale, tangible example. For instance, if it’s an optimization algorithm trying to solve a problem, I might do one country or one region or one product and say, “See, this is what we can do. I can give you this, but I need one or two people to solve it.” In my experience, showcasing results consistently helped me secure resources. If people can see a proof of concept – not just a talk or a presentation – then they’re happy to support.

I’ve often struggled to work with technology teams, especially engineering folks, who would happily collect data but were not concerned about how it would be used. This was particularly prevalent about 15 years ago when working in clusters was a popular approach. So, when I started running my own data teams, I insisted that they were responsible for the complete data lifecycle – collecting, storing, processing, and using data – I built full stack data teams. This creates a unified, cooperative environment, reducing friction and fostering team spirit.

LISA: Where have you had the most impact?

KSHIRA: I’d say it’s in making the entire business more efficient. I take the financial P&L approach to solving a problem, giving each team their own line item to answer the challenge. This helps each team understand how their work directly influences the business’s bottom line, keeping them motivated and the company profitable.

For example, if one team’s job is only looking at, say, labor cost, this team will just optimize labor cost and they’ll try to build all the relevant algorithms. If the other team’s job is just to improve the basket size, they will just focus on that. They can say, “I’m not just building a faceless algorithm, my work actually moves this line and therefore we are profitable by this much and we are very proud of it.”

I've tried iterations with one person doing multiple roles; they can do the data engineering part, be a data scientist, and build a dashboard. But it's very rare to find people like that and very hard to scale. Instead, we build a team with different skills, a mix of data engineers, data scientists, and business analysts. They can interface with different parts of the business but solve the same problem. That's a full stack team and it definitely works better.

LISA: Have you had to abandon a project for any reason?

KSHIRA: Yes, in the financial services sector, we began developing algorithms to predict credit scores before applications. It was an incredibly accurate tool, but we stopped midway due to ethical concerns and because we saw some people were interested in what we considered to be potential misuse of it.

There were a lot of passionate social justice champions on the team, and they will not build or work on something which they feel could be misused.

LISA: What advice would you give to companies on achieving value from data and AI?

KSHIRA: Start with the business problem you need to solve, not the data or technology. I've worked with a lot of leaders who get really excited about a new technology solution, but it may not align with the business strategy or contribute to efficiency. Blindly following technology trends isn't the answer. The key is to identify the business problem first and then work backward, exploring data-driven solutions. Otherwise, even substantial investments may result in unused technology.

LISA: Is there anything that comes to mind that you would do differently now?

KSHIRA: About 10 years ago, when we were big advocates for open source, we had to roll out a data dashboard project, but we focused too much on the destination and not enough on the journey. I learned that whether it's your own team or the enterprise, everyone wants to be on the journey. So now, I prioritize keeping everyone involved and informed, balancing pace with inclusivity, to ensure a shared sense of accomplishment and understanding. Knowing that people want to feel that their voice is heard has fundamentally changed how I operate.

LISA: Has there been a failure that's gone on to equip you for success?

KSHIRA: Yes, we once over-engineered a marketing attribution model which led to ridiculous results (10,000 ROI on \$15) because we got lost in the details. If you start going too deep and too granular, people don't find value in it. It was a great lesson in knowing when to stop. It's about balancing mathematical accuracy with practical usefulness. People need clear, actionable insights, not overly complex reports.

LISA: What excites you the most about data and AI now?

KSHIRA: I love getting to be part of what I call moving line items on a P&L. And for this to happen, I believe there needs to be a close partnership between business and data teams in the future. We can't just build tech solutions and walk away; instead, by working together we can significantly drive business growth.

Secondly, the potential to use data science and AI to streamline work and uncover inefficiencies is incredibly exciting. By minimizing repetitive tasks, we can focus on solving more complex and interesting problems.

LISA: What's your top tip for someone embarking on their data and AI journey today?

KSHIRA: One thing I tell the C-suite is to make room at the decision-making table for your data folks and make sure that the people who sit at the table are also leading those data teams. View them as value generators. Hold them accountable for generating and delivering tangible value for the organization.

Key Takeaways

- **Focus on the people, not just the tech:** Invest in working alongside people to understand their workflows in detail in order to demonstrate how data and AI can augment, not replace, high value processes. It's important to understand that deriving value from data and AI is not merely about handling numbers or algorithms.
- **Build-measure-learn cycle:** Take a financial P&L lens to solve a problem, by assigning each team their own line item to answer the challenge. Adopt a "build-measure-learn" cycle to keep efforts aligned – building the solution, tracking its impact, and constant improvement.
- **Full stack teams:** Deploy a full-stack team approach with different skills, including data engineers, data scientists, and analysts collaborating on the entire data lifecycle. Full-stack teams can interface with different parts of the business but solve the same problem.
- **Building usable AI solutions:** Start with the business problem to be solved, not the data or technology. Ensure AI solutions directly address specific business needs, to avoid investments in solutions that are unused. Ensure the output meets a specific business need to avoid it becoming unused despite substantial investment. Prioritize keeping everyone involved and informed, balancing pace with inclusivity, to ensure a shared sense of accomplishment and understanding.

- **Balance rigor with practical human logic:** Avoid over-engineering a model, balancing mathematical rigor with practical human logic to solve the problem at hand. Understand the pros and cons of each method or algorithm.
- **Integrate AI expertise in the C-suite:** Establish a close partnership between business, and data and AI teams. Integrate AI expertise into the C-suite, holding them accountable for delivering tangible value to the organization. View data and AI expertise as a value generator.



23

Joe Romata

Global Head of Customer Experience at a Multinational Energy Company

No matter how advanced the technology solutions are, their value is lost if they are not understood or used. The primary focus should always be on the perspective of the end user. Understanding their needs and interests allows us to engage in meaningful discussions about value, rather than simply pushing our solutions out. The key question to keep in mind is: “What’s in it for them?” This approach ensures that our efforts are targeted and effective.

LinkedIn: joeromata

JOE ROMATA

Joe Romata is the Global Head of Customer Experience at a multinational energy company. He has driven customer-centric digital transformations at a variety of companies, from blue-chip multinationals to tech startups.

LISA: Tell us about your career journey to date.

JOE: Following my time as a Deloitte consultant, I worked at a MarTech startup. Now I'm working for a multinational energy company as the Global Head of Customer Experience, looking at customer insights and analytics, channel strategy, and all the customer experience capability programs across the group.

Since 2008, I worked across a broad set of business challenges including operating models, business process improvements, digital strategy, customer strategy, customer journeys, personas, customer insights, and analytics. Then I started looking at the data to get the right insights for the business, reporting, exploration, and root cause analysis.

When I joined MarTech, the company was looking at using AI, back when it was a machine learning capability rather than generative AI. We were figuring out how to create marketing content and how to test it across channels to drive conversion. I worked on different campaigns with marketing departments from Google to Hotel.com. I was a solutions consultant as well as an account manager setting up scalable processes.

After the MarTech startup, I was headhunted into my current role in 2018 working for a multinational energy company to drive the customer-centricity agenda across the organization. Until then, the company had been quite a traditional product-led organization, not really customer-led. My initial task was to transform their Voice of The Customer program and generate the right insight.

When I joined, customer experience was just about one score of satisfaction, with no insights, and nobody knew what the improvement areas or key pain points were. The data was all over the place. I successfully transferred my role to leverage the latest technology. So, now we've got lagging indicators, which give us more information about what has happened, as well as lagging and

JOE ROMATA

predictive insights, which are improving all the time. These help product teams drive enhancements across the customer platforms. We're looking at employee experience metrics as well. The next step is linking all these customer and employee measurements to business performance. Hopefully, we'll get there!

In addition to that, I led the work on building our internal customer experience capabilities. What that means is establishing a common language, common training, and for everyone to understand the customer journeys. By encouraging a deeper appreciation of the customer experience, we can design processes that truly prioritize customer impact.

Most recently, I defined our omnichannel strategy and looked at the transformation program for next year. We've got limited resources and efforts – so, where do we want to focus? What are the key use cases that will move the needle and prove the value, and then scale?

It's all based around data and data-driven values. My role in the energy sector is unique. Unlike most positions that focus on a single business line, I have a broad scope, similar to our downstream director, working across all our B2B and B2C businesses.

LISA: Can you give an example of a specific problem you've solved with data and AI?

JOE: I pioneered an omnichannel customer experience measurement aiming to save \$40 million in costs, so that's a nice number. The challenge was that product teams relied on anecdotal discussions rather than having systematic insights. The solution I developed was analyzing data from our strategic B2B and B2C platforms, like the energy company's equivalent of Amazon (eCommerce, CVPS service platforms – customer value propositions – as a servicing platform) to identify customer pain points. These were across product ranges – from traditional fuels to new energy solutions.

To do this wasn't straightforward. We had to ensure a robust data foundation which starts with identifying what and where elements of customer data link to various points in the customer journey. Basic data, like order history, was often missing, hindering our ability to gain an accurate view of the customer. We have just started merging data from siloed systems and improving the data quality, then we can access meaningful customer insights.

LISA: When you started working at the energy company, how was data organized and how did you change it?

JOE: Up until 2019, the business's attitude was very much, "We make a product and they will buy it," so the business wasn't thinking about what customers really need. Then, there was the realization that when competing with startups and in some sectors, such as home electrics and home electricity charging, which were becoming saturated, you actually have to compete differently.

To create an organizational cultural shift, we needed to pay more attention to measuring outcomes and customer success. Additionally, we needed customer insights to understand our different customers' needs and how to meet them. There are different sources for these insights, including market research, account managers, and frontline colleagues; however, they were not systematically measured.

When I joined the company, the way we measured customer experience was with one metric of "how satisfied are you?" or "would you recommend the company?" There were no material insights. There were some verbatim questions, but without technology that enables you to analyze them, there was no one really to look at them. We had gathered feedback but nothing was done with it, so none of those insights were used.

Colleagues [at the energy company] wanted to build something and grow the business and work out what to focus on to really move the needle. We needed baseline data about the customer so we could actually measure and report effectively. We had some anecdotal understanding of where pain points might be, so we started measuring where we could make fast improvements.

We didn't have refined customer journey maps back then, so we had to rely on the key moments. For one of our strategic platforms, our customer insight was dispersed, so we had to logically think about where we should start. Is it the login process, or is it the ordering process, is it the invoicing process? We teased out the key pain points across the platform informed by customer feedback and deployed text analytics capability to mine customer comments verbatim. This gave us context to the score.

LISA: How did you make a plan and what were the key things you needed to deliver initial value to the business?

JOE: One of the most important assessments we performed was to identify the program's value proposition for our internal business users and in turn our customers. To achieve this, we interviewed our internal business users to understand their challenges and educated them on the value of the program in addressing those challenges.

Next, to create a customer profile, we started by breaking down the program's key components of the customer data and understanding what value could be gotten out of them without asking another 50 questions. Then we tagged the data to ensure measurement points throughout the customer journey.

At each stage, we had to engage with different teams, data management, analytics, data privacy, and GDPR compliance to work out what we wanted to use for the system. To establish the right data foundations, we needed to get an appropriate vendor to support and develop the right roadmap for the next two, three, four, five years.

We needed a culture change to deploy new ways of working. We embedded insights into different everyday forums and discussions. We made a shift from a waterfall approach to more iterative cycles informed by these insights.

JOE ROMATA

LISA: Have you started to work with generative AI?

JOE: We are integrating more generative AI into our channel management. In fact, it's forecast to reduce Average Handle Time (AHT) by 20%.

LISA: What strategies did you use to overcome resistance and get everyone behind you?

JOE: The number one thing that worked really well, which I learned from consulting, is to understand the business end user's angle. What's in it for them? What I mean by that is that each business user is interested in their function and specific challenges.

For the product teams, we created a dashboard they could use internally to understand what they build and how that impacts customers. Quotes and feedback from customers helped bring insights and customer personas to life, and helped the team understand all the interlinks and potential benefits for product improvements.

For the sales team, once they understood how we could resolve the customer pain points, they had more time and success in sales conversion.

The other key to success is strong leadership support. They understood the value, and, once I got their buy-in I had the support to go ahead. This reinforced the approach to leveraging insights, which are now readily and intuitively accessible for users as needed.

LISA: How did you align disparate data sources?

JOE: We used tools to pool data from different sources. The data didn't sit in a warehouse but was set in different systems. We had SAP CRM and Salesforce. We had the digital platforms or customer platforms, and we pulled the data together, manipulated it, and then provided it in the most appropriate format for our purposes. We had a data lake program in process, but we couldn't wait for completion before commencing this work.

Key Takeaways

- **Effective customer experience measurement:** Ensure strong data foundations with the right data points readily available to connect customer interactions across all touchpoints. Optimize data quality and accessibility to gain a comprehensive understanding of customers. Breakdown silos, merging data across different departments in collaboration with data management, analytics, data privacy, and GDPR compliance teams.
- **Culture of customer-centric measurement:** Augment cultural shift to effectively measure customer experience including customer satisfaction, capturing efficiency-driving insights, and embedding them into daily work and decision-making forums. Include customer feedback, which is often overlooked. Leverage insights to not only improve customer experience but also reduce operational costs and inform the strategic implementation of AI. Adopt iterative work cycles that are informed by customer data, allowing for continuous refinement and adaptation.
- **Building buy-in and momentum:** Take a multi-pronged strategic approach. Understand the business end user's perspective and their challenges. Empower product teams with data that helps them understand the impact of their work on customers and potential improvements. Increase sales conversion by providing data to the sales teams, so they can understand and address customer pain points. Win leadership support by demonstrating the value of the work.



24

Tomasz Ullman

Former Global Head of Data Science and Strategy at Ford Pro

AI and machine learning are based on probabilities. There's no certainty. We are building probabilistic models. We underestimated the challenge that this would pose. You need to make sure that C-level executives understand both the power and limitations that data science, AI, and machine learning can offer.

LinkedIn: [tullman](#)

TOMASZ ULLMAN

Tomasz Ullman is the Commercial Director at RAC. At the time of this interview, he was responsible for Global Data Science and Strategy at Ford Pro Tech Data and Analytics. Prior to this, he was the Senior Analytics Manager at Lloyds Bank and a Manager at Deloitte. He has extensive experience in leading large-scale commercial and data science initiatives.

LISA: Can you tell us about your career up to now?

TOMASZ: Currently, I'm Head of Data Science at Ford Pro, a business unit within Ford Motor Company that focuses on commercial customers. That's everyone from self-employed plumbers to massive global fleets for Amazon. It's vehicles and, as we move to battery electric vehicles (BEVs), it's also increasingly charging. Software and what we call productivity tools are also becoming a big part of the experience of vehicle usage. We're creating the productivity tools around the vehicle that help businesses focus on what's really important to them. So, it's service and parts financing, using data science and analytics solutions for those different verticals. Some of it is to help executives and leaders of those verticals to make informed decisions and be smarter, while some of our initiatives focus on building data-driven products that we can then offer to customers.

Prior to my current role, I was the Smart Mobility Analytics Lead for Ford in Europe, focusing on future mobility solutions. Data was a big part of this as we were designing all kinds of solutions for anything to do with mobility, so it included solutions for cities and insurance companies as well as vehicles and public transport, as well as micro-mobility like scooters and e-bicycles.

Before that, I spent a couple of years at Lloyd's Banking Group in their marketing analytics function, leading a brand analytics team. And before that, I was at Monitor Deloitte.

LISA: Could you give an example of how you navigated using data within one of the companies you worked for?

TOMASZ: When I was part of Ford Mobility, it operated a bit like a consultancy for cities. Public transport organizations like Transport for London are often public institutions, so they can be slow-moving, which presents its own set of challenges.

This particular piece of work we did in 2019 was in London where there's lots of legislation around congestion charges and restrictions on large goods vehicle lorries, which are seven and a half tons or more. These massive lorries are only allowed to go on predefined major roads through all 32 London boroughs unless they have specific permission to use smaller roads. If they don't comply, there should be a penalty, but in 2019, there was no enforcement mechanism for this; so, unless the police happened to be driving right behind them, there were no penalties issued.

One borough came to us and asked, "Could you help us design, build, and implement a solution to help us enforce this legislation?" The things that we had at our disposal were data and data science. They wanted to invest in a network of cameras that would help track vehicles' movements. They wanted analytics and a data science engine to be integrated with their ticket-issuing system to allow them to operationalize it.

It couldn't be enforced with a simple one-spot camera, like the Ultra Low Emission Zone (ULEZ) or speeding cameras, because this is about catching a vehicle on one of the hundreds of thousands of small roads. But you can't put a camera on every intersection and every road, which is why there was so much leakage in the enforcement – a problem not unique to this borough.

First, we needed to assess the size of the problem. So, we purchased a dataset of traffic flow from companies that capture this data from plug-in car devices. This data included one representative month in the year that recorded flows of vehicles fitted with Pit devices or embedded modems.

In that one month there were half a million trips and fifteen million data points; these were basically breadcrumbs. We now knew where they started, where they finished, what they did, and when, with timestamps. We could then map it to see the volume of traffic flow through the different roads and how many of those trips were violating the law.

We knew exactly where people were going and how many trips originated and ended outside and inside the borough. The city then hired a civil engineering traffic company that installed vehicle counters for a week to recalibrate the volumes we had gotten from our months' worth of data. Then we could look at those volumes and traffic patterns to assess the opportunity. How many vehicles are driving only on permitted roads? How many are deviating from those routes and violating the law? What can we capture? Then we thought about what the enforcement solution might look like, came up with two or three, and settled on one approach with two elements.

The first part was a sequencing model that looked at the sequence of passing cameras and all the possible combinations. We started with a network of the 50 cameras our budget stretched to, then looked at placing them at major intersections, and what would be allowed and not allowed in the sequence of passing the cameras. If you observe a vehicle passing camera A, and then the next time you see it, it's passing camera C without passing camera B, then you know they violated the law because the only permitted route is to go past camera B before C.

We had a massive spreadsheet with all the different possible combinations, like a correlation matrix on all the cameras and different rules. Then we took the actual data from the trips and mapped it onto the matrix, so we could see how many trips would have been picked up by different combinations of sequencing cameras. This enabled us to draw a detection curve. The first few cameras you install provide really high value and, from that, we could build more sensitivity into the curve. By adding additional cameras, what incremental coverage do you add against decreasing marginal returns? Then, by adding in the costs of the camera, ongoing maintenance, the number of tickets you can issue, and the expected revenue, you can calculate the breakeven point

and provide a recommendation on the optimal number of cameras. That was phase one: detection. Obviously, when you build something like this, you're going to have a number of trips that are incorrectly classified or maybe there was an obstruction. The cameras aren't perfect; you might miss something, or someone might appeal their ticket and so on. That's one of the challenges of working with leaders who don't really understand what data can offer them.

Data science and AI are struggling with the fact that London and other cities can issue tickets but they need to provide clear evidence of the violation. If you get a speeding ticket, then you get a picture of your car and a timestamp and the speed you were going. In our case, the evidence is the absence of you showing up on the camera. So, the whole logic of evidence is completely upside down. The city was scratching their heads, thinking, "How are we going to evidence this and be compliant?" This was a nationwide problem.

The second part of our modeling was, if you have the correct sequence where, you expect to see a vehicle going between camera A and camera B, that vehicle may still deviate between those two cameras into a residential road. So, our model assumes there's a time element to it, but how much time, and can you account for whether it's rush hour or not? The time element was the second area of enhancement.

The problem again with this is that the basis for detection of the violation is probabilistic. Are we right, probabilistically assuming or statistically assuming how long it should take? We look at the curve at this point – say twelve o'clock on Tuesday – this is how much traffic there is, this is the congestion, this is the curve so, if you fall outside of that normal distribution, then you know with a certain degree of confidence you're likely to have violated the ban. This provided an additional complication of how to deal with evidence that's not binary; it's not using the evidence, it's based on probabilities.

The work was received as a clever way of looking at it. We could show exactly what detection rates are, we could build a compelling business case, and so on. But the problem started with the practical side of things. And then the main stakeholder we were working with, who was our biggest supporter, moved to a different role. So, in the end, the work didn't get implemented.

When I look back on it, I think the leaders didn't understand that AI and machine learning are based on probabilities. Unless you were able to use something like camera density, to do perfect vision-based tracking, there's no certainty. City officials are a lot like the automotive industry; they are used to dealing with engineering facts, which are measured with devices. It's very binary, so the way people make decisions is very black or white. Yet, the new tools of data science – you don't need to go to machine learning and AI, just simple statistics or econometrics – are all based on probabilities. This made the city officials uncomfortable.

It also raised all sorts of questions, including legal ones, as officials anticipated people pushing back on tickets, saying, "You have no evidence to prove this." So, then you start to look at all these other operational costs, such as how do you deal with these complaints? It became a more complicated problem for the city.

LISA: What would you do differently next time?

TOMASZ: We should have had a better relationship with more senior decision-makers so we could have discussed these issues and scenarios before building the models, designing the solution then building it, and so on.

We underestimated the challenge that building probabilistic models, which have an inherent level of uncertainty, would pose. It's almost a philosophical approach to what you accept as fact and what you accept as evidence.

This was just one example. We did a lot of work with London and other cities where the data science solutions we would propose had that element of probability.

LISA: Can you share an example from your time at Ford where you successfully built, scaled, and implemented an initiative?

TOMASZ: Certainly, a significant challenge for automakers today is the transition from internal combustion engines to BEVs. Technology in this space is moving faster than consumers' understanding of what it offers, so they, and businesses,

often struggle to grasp the implications. Businesses, particularly those with large fleets, are making large investments so, it's a big decision for them. In addition, some companies, especially delivery companies, have environmental targets and want to reduce their emissions. We wanted to provide them with the necessary insights to make informed decisions about the pros and cons of switching.

So, we developed an online calculator tool on Ford Pro UK. For existing customers sharing drive data with us, the tool performs real-time analytics to compare the performance and costs of internal combustion engines with BEVs. Questions like, "Is electric a suitable replacement?", "What range do you need?", "What temperature are you operating in (air conditioning puts pressure on energy consumption)?", "Can you charge overnight?", and "What are the operating and maintenance costs?" These are all answered by the tool, tailoring the advice to each individual business.

You can go to the Ford Pro UK website where the first version of that tool is up and running. If you're an existing fleet that already drives vehicles with connected modems and you're already sharing data with us, then you can log in, and the analytics engine looks at the real drive data of that fleet based on the internal combustion engine behavior. You can translate it to energy, and from that you can translate it into what it would mean for a battery alternative. It's a like-for-like comparison. What would it look like if you were to drive a battery electric vehicle today?

The trigger to build this tool is aligned with Ford Pro's strategic goals to boost BEV sales in line with annual targets up to 2030. Achieving these targets will naturally cascade to other developments, like more charging stations and more advanced telematic solutions to track. In addition, there are broader implications for upskilling dealers to service BEVs, diversifying supply chains, and more. All of these things are built into our growth plans.

LISA: What challenges did you face in building the tool?

TOMASZ: There were two main challenges: technical and legal. Consent was crucial because we could only use data from registered customers who explicitly granted permission, which led to a lot of work on data segregation and legal compliance because the data could only be used for this initiative.

On the technical side, we had a trade-off between model accuracy and applicability. The tool is highly accurate because it takes a specific set of data on the specific fleet and does a like-for-like comparison, but it's restricted to existing Ford fleet customers with relatively new vehicles that have connected modems. It doesn't work for older models or non-Ford vehicles, which limits its broader application. So, right now this tool is good for upselling or cross-selling for existing customers but limited beyond that. Future phases aim to simplify this model and make the tool more inclusive, allowing customers to fill in the details themselves. We can then provide some high-level benchmarking that's not as accurate but gives a broader audience a flavor of the results. Right now, our Ford Pro Telematics can integrate non-Ford vehicles through API integration or with a non-OEM plug-in device. We're considering extending this feature to the new tool as well, making it more brand agnostic, so we could potentially assess a non-Ford engine fleet to provide similar recommendations.

LISA: How do you measure the impact of a tool like this?

TOMASZ: We can measure that people who went through the process were more likely to switch their fleet to battery electric compared to those who didn't. We have sales teams pushing this product and making their pitches. We've seen a good uplift so far but it's too soon after launch to say more right now.

LISA: What do you wish you'd known about building a tool like this?

TOMASZ: I wish I'd understood the depth of the knowledge gap among the fleets we target. Because they've never driven a battery vehicle, they often don't even know the right questions to ask about transitioning to battery electric

vehicles. While customer-centric design is essential, it's not entirely effective here because the customers themselves are unaware of their future challenges. We say, "Go and talk to customers to understand their pain points," but in this case, they don't know what challenges they're going to run into. So, I'd say, "Sure, involve your customers but don't expect them to give you all the answers." It's like Steve Jobs designing the iPod even though it wasn't solving a problem, and nobody was asking for it; he knew it would be appreciated. In the same way, people don't know what to expect from BEVs so you have to guide them – innovation needs to anticipate needs.

Also, it's crucial to have executive buy-in, particularly from C-level leaders who can understand both the business value and the technical aspects of data science, AI, and machine learning. A data-savvy CFO or CMO can better articulate the tool's value and benefits, making it easier to secure support and resources for development. You have much more success if you have a data-savvy CFO or CMO who can speak a language of value and benefits and bottom line while understanding enough of the technology and data.

LISA: What are you most excited about with data now?

TOMASZ: Traditionally, automotive companies only thought in terms of vehicles where cars were simply viewed as a means to move from point A to B. Now, with data, software, and electrification, vehicles are evolving into multifunctional productivity tools – serving as workplaces, entertainment centers, and more. The emergence of AI further amplifies these possibilities, revolutionizing the entire mobility and transportation landscape.

The whole area of mobility and transportation is completely changing.

How all this intersects with urban development and new technologies such as the Internet of Things, is very inspiring. These new players and dynamics are all enablers of an exciting new future. The mobility ecosystem is increasingly being driven by data which is changing how we view transportation and also impacting how we live and interact, and that's super exciting to me.

Key Takeaways

- **Probabilistic decision-making:** Invest in educating stakeholders who are accustomed to dealing with binary facts on the full scope implications of probabilistic decision-making in AI. Run scenarios of potential outcomes to demonstrate how a model might impact decision-making. Anticipate potential legal and operational challenges that probabilistic models can introduce, and plan mitigation strategies beforehand.
- **Future-proof executive support:** Build a broad executive understanding of the financial impact and proposed AI solutions to gain buy-in. Recognize that the development and implementation of AI solutions often require time and investment. Secure this wide support to mitigate the disruption risk of key stakeholders moving to different roles.
- **Data consent and legal compliance:** Understand the industry's specific regulatory and enforcement challenges. Secure explicit consent from customers to use their data for a specific initiative. Ensure strict segregation from other datasets and that it is only used for the specified initiative.
- **Technical trade-offs:** Consider the trade-off between model accuracy and applicability. While a model might be highly accurate for a specific set of data, it may not have a broader application.



25

Oz Krakowski

**Chief Business Development Officer
at Deepdub**

There's this huge push for self-identity – it's like everyone wants things tailor-made for their own culture, preferences, and personal identity. This isn't just about movies or TV shows; it's everywhere. And, when you project this trend into the future, it's clear that making localization easy and accessible is going to be a game-changer.

LinkedIn: ozkrakowski

OZ KRAKOWSKI

Oz Krakowski has a solid educational background in electrical engineering and holds a Master's in Business Administration (MBA) from the SMU Cox School of Business in Dallas. With 16 years of experience in the US working with international clients including Texas Instruments, he seamlessly combines a unique blend of business acumen with technical knowledge.

Oz's entrepreneurial brothers founded Deepdub in 2019 in Israel. Oz joined this venture which aims to transcend language barriers and cultural gaps in audiovisual content. Integrating AI, deep learning algorithms, and automated workflows, Deepdub meets the burgeoning demand for localization in the media industry by making content universally accessible and culturally resonant.

LISA: How did you and your brothers decide to focus on the media industry?

OZ: The advent of digital streaming services and the effects of the COVID-19 pandemic revealed a significant need for accessible localization, which Deepdub aimed to fill. We offer both a managed service and a SaaS platform, catering to a range of customers from B2B to enterprise clients.

The whole issue of localization is much more complex when you look beyond entertainment. Remember when everyone used to talk about the world being a global village? The idea was that we're all connected and can hop across the globe in a heartbeat. But nowadays, there's this huge push for self-identity – it's like everyone wants things tailor-made for their own culture, preferences, and personal identity. This isn't just about movies or TV shows; it's everywhere. And, when you project this trend into the future, it's clear that making localization easy and accessible is going to be a game-changer.

Why shouldn't people in Southeast Asia, for example, enjoy content in their own language? The only reason they haven't until now, is because it's been too pricey or just unavailable. There's been no proper setup for it, so they've had to make do with subtitles. But, for example, when literacy rates are low, subtitles aren't going to cut it.

So, what we're looking at is dubbing stepping into the spotlight. It's going to become way more prevalent. Sure, entertainment is a great way to test out the concept and get the ball rolling, but the potential market for localized content is massive beyond that. Think about it: there's a whole world of audiovisual material out there, and people are more than ready to consume it in their own language when it's available. That's the future I see, and it's a pretty exciting one.

LISA: Can you explain how Deepdub provides solutions for entertainment that might otherwise not get translated or dubbed (localized, as you term it)?

OZ: Many of our clients wouldn't localize their content if not for our technology, for various reasons. Even those who do localize at scale often come to us to expand their current services. We can pretty much open a new territory because our technology dramatically lowers the barrier to entry compared to traditional localization processes, making content that wouldn't be economically viable to localize otherwise feasible. It's simple economics; if you lower the cost of investment, the greater the chances that you're going to see the return on it. And the idea for us is not only that we're making it more economically feasible, but we're also dramatically reducing the turnaround time, which means you can start monetizing it faster. So, instead of being two years in the future, it's going to be three months. Here's a good example: we worked on a Danish show called *Follow the Money* – a mix of Danish and Arabic, which, by the way, adds a whole new layer of complexity when you're tackling localization because it's not just about the language, it's also about the cultural nuances, right? The show had three seasons, each with 10 episodes, and each episode runs for, I'd say, about 57 to 58 minutes, which totals over 1,700 minutes of content to dub. Traditionally, dubbing a single season of 10 episodes could take 8 to 14 weeks, which includes everything from the moment you get the materials until you hand over the final product to the client. And it is a convoluted process with many additional steps – setting up, casting, project management – way beyond just producing the dubbed content.

But in our dubbing service, we processed all 30 episodes in just 8 weeks; that's a 60% time saving and cost saving of at least 30%. It's a huge deal, especially when you consider that some of the big studios would typically shell out up to 50% more for traditional dubbing services.

And we're not just talking dubbing here. We deliver the full package, from end to end; no need for extra investments for those additional, often pricey, side products. From the materials we receive to the finished, localized product, we handle it all. So, not only do we save time and money, but we also simplify everything by delivering a ready-to-go, localized show.

LISA: Can you talk us through the technology stack behind your solution?

OZ: We've developed two main business units: our subscription-based service and the managed services, which are like two sides of the same coin, technologically speaking. Our subscription-based service offers end-to-end localization at scale on a single platform. The subscribed users drive the work, and the platform enables them to control every step of the localization process while they can determine how much or how little they wish to intervene in the automated processes. Our proprietary technology is the first of its kind with emotion-based text-to-speech (eTTS™) and speech-to-speech capabilities enabling the user to have control over the outcome.

Deepdub's localization workflows

Hybrid	Automated (Deepdub GO)
High quality, human-led, AI-assisted Designed to support dubbing for feature films and episodic scripted content	Low cost, super fast, automated, AI-based, human-assisted Designed to support dubbing for factual and non-scripted content Designed to support FAST channels

Figure 25.1 – Deepdub.ai Localization Workflows

For our managed services, Deepdub aims to bridge the language barrier and cultural gap of entertainment experiences for international audiences. We provide high-quality seamless localization of entertainment content using deep learning and AI algorithms. This service plugs into the post-production process of content owners and takes complete ownership of all their localization needs. It uses a combination of our unique proprietary technology, an in-house expert production team, and the relevant certifications.

The managed service is super versatile, letting us handle a wide variety of content and nail down a high-quality finish. On the other hand, the SaaS platform is more specialized, dialed into particular types of content and workflows. It's perfect for when you need good-enough quality without all the bells and whistles.

We employ in-house-developed generative AI models for the entire localization workflow, which is broken down into transcription, translation, adaptation, voice generation, and mixing. The dubbing process starts with ingesting the video content. If we don't have a script, we transcribe the audio first, turning speech into text. Once we've got the transcript, we move on to translation, which is this neat dance between AI and human smarts. The exact mix depends on what we're working on and what result we hope to achieve.

After translation, there's adaptation. That's where the human touch really shines. Different content needs different tweaks, right? Like a joke in Danish has got to land just as well in English. Then comes voice generation. Instead of booking studio time and having voice actors do their thing, we're creating voices with computers. It's pretty wild. We can start from plain text or even morph existing speech into a new language.

The last piece of the puzzle is the mix. That's where we blend the new dialogue with all the original background tunes and sound effects. It's like recreating the original vibe of the show or movie. If we're working on something simpler, like a training video, we can skip the fancy mixing. But for entertainment? It's got to be spot-on, or else it'll just feel off.

What's really slick about our platform is how we've streamlined the process. Imagine one person handling transcription, translation, and even voice generation – all in one go. And the mixing? It's automated. So, you press a button and bam, you've got a fully dubbed video ready for YouTube or whatever.

For more intricate stuff like *Follow the Money*, we go all in, fine-tuning every detail. The AI might be smart, but it's not cracking jokes yet, you know? That's why we still loop in native speakers to make sure everything's on point.

One of the big perks of our platform is how we let humans step in at any stage to make sure the tech output hits the mark. Because let's face it, when you're sinking into your seat at the cinema, you're not going to settle for less than perfect. And we wouldn't want to be the ones getting roasted on Twitter after the premiere, right?

LISA: How do you see customer processes changing with the introduction of your technology?

OZ: The integration of our technology is really simplifying customer operations, especially where traditional localization has needed a lot of effort and coordination with external parties. Localization has been slow to evolve due to a lack of technological advancements but we're starting to see a shift. As our AI becomes more accurate and requires less human curation, we anticipate a future where live dubbing and seamless localization are commonplace. As the need for extensive human intervention diminishes, I foresee large operational teams will shrink but this doesn't mean job losses; rather, it's a transition. The workforce will migrate toward different, and often new, roles that are emerging alongside AI developments. The effort required for tasks such as localizing a TV show will decrease dramatically.

LISA: What changes do you foresee in studios' production processes and the future applications for your technology?

OZ: Looking to the future, I think the key term is "better." By that, I mean more accurate and requiring less curation. Curation won't disappear; it'll always be necessary, particularly in entertainment, to meet audience expectations. However, as our technology evolves, becoming more accurate and capable of handling a wider range of idioms, the necessity for intensive curation to achieve quality output will decrease.

At the moment, the speed of content processing is limited by the need for curation at multiple stages, including translation and voice generation. Although the quality of our output is already impressive, there are occasional glitches that we are aiming to eliminate. This isn't a distant dream; it's a reality that's almost at our fingertips.

We're on the verge of being able to offer glitch-free voice generation. Coupled with significant improvements in translation, we're moving toward live dubbing capabilities. Some startups are already experimenting with this, albeit with poor quality. But we're focusing on quality. Imagine live events broadcast in multiple languages simultaneously – that's where we're headed.

The trajectory of our technology will allow for live dubbing in original voices across multiple languages. Video lip-synch will eventually enable actors to perform in one language with their lips moving accurately in another. As this technology is integrated into workflows, it will be possible for me to speak to you in Hebrew and you hear me in English, seamlessly. This is the transformative potential of our technology, and it's an exciting prospect.

LISA: Can you tell us more about your use of AI in the post-production services you offer studios?

OZ: One of the most innovative features we're developing with studios is voice cloning, which is particularly useful during the ADR (Automated Dialogue Replacement) process for film and TV production. For example, after shooting

an entire TV series in one month, it might go live two years later, during which time it undergoes multiple production processes, including additional viewings and testing, which often require dialogue changes.

Traditionally, this means calling back actors to re-record their lines, which is expensive and logistically difficult, especially if they're involved in other productions. Voice cloning allows us to make necessary dialogue changes without physically bringing the actor back into the studio. This can be a cost-effective solution, reducing both operational burdens and expenses for studios. Moreover, consider a situation where the actor was a child during the original filming and their voice has changed by the time pickups (additional filming) are needed. Voice cloning can replicate the original voice exactly as it was, which is invaluable.

Although it's a niche market – since it mainly serves high-end film and TV shows, of which there are a limited number produced each year – this technology is incredibly valuable. It's not just about the 500 or so films produced by big studios annually; it's also about the thousands of hours of content that need localization.

LISA: Are there any key messages you'd like to share that could benefit CEOs and board members?

OZ: I think it's crucial to recognize that the landscape of AI is shifting rapidly, and inaction is not a safe strategy. The reluctance to embrace AI due to uncertainty can be detrimental. From my conversations over the past six months, a common theme is the hesitancy to adopt AI, often stemming from a wait-and-see attitude. However, in our fast-evolving world, waiting even a few years to see where AI is heading is excessively risky.

C-suite executives and influential leaders begin to see the value of AI once they engage with it, particularly when the conversation turns to the financial implications. The potential impact on the bottom line is significant, and understanding this is critical for senior executives whose primary concern is economic outcomes.

Unlike previous technological revolutions that spanned years, the pace of AI development is happening quarterly. For example, recent advancements allow AI to analyze images, and we're making leaps every quarter. If you're not on board, you're not just two years behind – you're in technology dog years, where every quarter equates to several years of progress.

Key Takeaways

- **Simplify customer operations with AI:** Consolidate repetitive tasks within high-value processes and leverage AI tools to automate tasks, handle inquiries efficiently, and ultimately enhance customer experiences.
- **Human-in-the-loop (HITL) approach:** Use AI for repetitive analytical tasks but maintain a human-AI collaboration for complex tasks and decision-making. This fusion approach yields the greatest high-quality outputs.
- **Explore subscription models for easier AI adoption:** Consider subscription-based and managed services models for a lower-risk entry point to explore AI. Experiment with different solutions without a significant upfront investment, and benefit from expert support from the service provider.
- **Transform customer connections:** Personalize customer experiences with AI. Incorporating AI into customer operations isn't just about efficiency – it's a fundamental shift that aligns with our evolving understanding of identity and culture.
- **Key messages for CEOs and board members:** Recognize that the landscape of AI is shifting rapidly and inaction is not a safe strategy. Overcome the common hesitancy to adopt AI, whether due to a wait-and-see attitude or uncertainty. Understand the transformative potential of AI and proactively integrate it into organization strategies.



26 CASE STUDY

LLMs and RAG Enable Hyper-Personalized Education for Healthcare Technicians

IN CONVERSATION WITH AMR AWADALLAH: Founder and CEO of Vectara, a generative AI software company with a mission to help enterprise customers find and leverage insights more effectively from their unstructured data assets.

LinkedIn: [awadallah](#)

Case Study in Summary

Case Study Classification

- Segment: B2B
- Industry: Technology, medical services
- Geography: North America
- Value creation levers: Product knowledge, training
- Time to impact: Less than 1 month
- Regulatory considerations: High

The Business Challenge

Vectara's customer, SonoSim, an ultrasound training and education platform, needed a scalable mechanism for end users (ultrasound technicians) to increase the speed at which they could be educated about how to use, calibrate, and optimize ultrasound machines and deploy advanced methods of detection and prevention. Legacy solutions relied on a traditional search query model that yielded a long list of results, requiring technicians to sift through large amounts of content to identify the most suitable response. The challenges were:

- Developing a method to navigate large volumes of technician-specific materials more intuitively and quickly, to get insightful, relevant answers.
- The ability to quickly draw inferences between highly technical terminology and serve grounded responses based on the training data.
- Ensuring a high degree of precision in answers while safeguarding proprietary content from being shared outside the company.

The Response

Amr and his team at Vectara have developed an end-to-end approach to help companies like SonoSim leverage powerful generative AI technologies using their own data. Their software solution helped SonoSim implement a more intuitive, natural-language-based capability enabling ultrasound professionals to rapidly query extensive materials to get specific answers to a variety of questions.

The work was done in two phases:

Phase 1: Defining the use case and knowledge discovery. SonoSim's extensive content repository spanned 9000+ pages of text (equivalent to 300+ textbooks), 5000 images, and 4000 video clips with continued rapid growth. They needed a mechanism to query this database on mobile devices and deliver precise teaching content tailored to each user's needs.

Phase 2: Implementing Vectara's platform and calibrating it to SonoSim's environment. Vectara's Retrieval Augmented Generation (RAG)-based generative AI cloud-based platform delivered a natural language search capability. RAG improves the accuracy of generative AI models by fetching facts from, and grounding the model on external knowledge sources. The conversational interface enabled interaction with real patient cases, delivering a hyper-personalized learning experience, custom course creation with continuous assessments, skill development, and mastery for medical professionals that overcame the previous financial, linguistic, and logistical barriers.

The solution had a significant impact, streamlining access to information while enhancing user engagement and satisfaction. Generative AI also harnessed previously underutilized data to improve the models and performance, while positioning SonoSim as a pioneer in state-of-the-art services to users in healthcare education.

Case Study in Detail

Amr Awadallh profile

As an innovator in the tech industry, Amr Awadallh has significantly influenced big data and AI in a career spanning over 25 years. As the CEO and co-founder of Vectara, he is at the forefront of transforming generative AI and conversational search using state-of-the-art technologies like deep neural networks, large language models, and natural language processing. His extensive background includes serving as the Vice President of Developer Relations for Google Cloud and eight years as VP of Product Intelligence Engineering at Yahoo. After this, he co-founded Cloudera in 2008, where he was the Global CTO. Cloudera went public in 2017 at a valuation of over \$5 billion.

Amr, a respected figure in the global tech community, holds a PhD in electrical engineering from Stanford University. His deep technical expertise, combined with a strategic business perspective, is a result of his tenure as VP of Product Intelligence. His desire to share insights further enhances his reputation and influence in the field.

The Business Challenge

AMR: SonoSim specializes in training technicians to expertly utilize, calibrate, and optimize ultrasound machines. As a life-saving technology, ultrasound imaging's effectiveness hinges on skilled detection and identification. Mastering this technology is complex, often requiring technicians and radiologists to seek frequent guidance. Inadequate expertise can compromise the critical use of ultrasound, as exemplified by procedures like guided internal jugular vein cannulation, where a mere 3mm distinction in imaging can be the difference between saving a life and causing a stroke. Such real-world applications underscore the importance of proficiency in various ultrasound modalities, including sonographic anatomy and hand-eye coordination to navigate a probe in 3D space.

LLMS AND RAG ENABLE HYPER-PERSONALIZED EDUCATION FOR HEALTHCARE TECHNICIANS

SonoSim's mission is to enhance medical professionals' ultrasound proficiency across various scenarios, offering a rich array of educational tools, resources, and services. The challenge lies in managing the vast, ever-growing content library to allow users – both practitioners and educators – to swiftly find specific information.

The value of bedside ultrasound is highly dependent upon the skills of the medical practitioner performing the ultrasound exam.

— Eric Savitsky, Founder and CEO of SonoSim

The company aims to streamline the education pathway, for more efficient machine usage and immediate problem-solving during procedures. However, the traditional search for information is time-consuming and cumbersome. Content generated by technicians, including manuals and FAQs, exacerbates this issue by increasing the bulk of material to navigate.

SonoSim, therefore wanted to implement a solution that accelerates the search for and retrieval of information, simplifies content creation, and promotes consistent practice for mastery. This effort is vital for integrating ultrasound into broader clinical assessments, which requires coordination with other busy healthcare professionals. The primary goal is to deliver effective teaching within resource, time, and budget constraints, fostering repeated learning opportunities to integrate ultrasound's role in clinical decision-making at a deeper level.

Context-dependent issues

Implementing generative AI technology in the medical space brings specific technical challenges that must be navigated with precision:

- **Reliability:** Above all, the AI's output must be dependable. The everyday reality of hallucinations or misleading AI-generated data, if unchecked, could be detrimental to clinical outcomes.

- **Patient privacy:** Safeguarding sensitive health information is non-negotiable. Any AI system must uphold stringent patient privacy standards to maintain trust and compliance with regulations.
- **Explainability:** The AI's decision-making process must be transparent, allowing healthcare professionals to understand and trust the logic behind AI-generated recommendations or analyses.
- **Bias mitigation:** AI systems must be designed to minimize inherent biases and ensure the integrity of the medical insights.

Optimizing reliability is the cornerstone for addressing the subsequent challenges in maintaining privacy, clarity, and fairness. Without ensuring that the AI's outputs are accurate and trustworthy, other goals remain unachievable. Therefore, reliability is not just a preliminary step but a continuous requirement that underpins the integrity and utility of generative AI in healthcare.

Approach

Phase 1: Defining the Use Case and Knowledge Discovery

AMR: The main friction for a business leader is confusion about generative AI applications. They say, "We know we need to use generative AI, but we're not sure what to use it for." Usually, the discussion happens very quickly with them. We start by asking, "What are some of the problems you have in your company?" Then we demonstrate how Vectara's solutions can help.

Some of them are clear, such as with SonoSim's need for information retrieval. It was very clear that the product could deliver answers deeply aligned with the user's search intent, far surpassing keyword-based solutions.

SonoSim's process involves scanning real patient ultrasound scans and aggregating the data (volumetric, 3D and 4D, motion sensing, and so on) into a centralized cloud platform to create a vast library of training content accessible via mobile devices. This cloud-based platform offers multiple modalities, assessments, and skill tracking for practitioners. The ever-growing repository is extensive, with

over 9,000 pages of text (equivalent to over 300 textbooks) more than 5,000 images, and upward of 4000 video clips.

SonoSim's use case was to implement the most efficient method for practitioners and educators to pinpoint the precise content they need, such as, "I'm looking to scan a large pericardial effusion," to deliver extremely precise, hyper-personalized teaching, tailored to the user's immediate needs.

To demonstrate the scale, context, and urgency of this challenge, today's medical knowledge is expanding at such a rapid rate that it doubles approximately every 70 days, compared to once every 50 years in 1950. So, SonoSim sought a solution that could deliver the correct educational content swiftly in this ever-expanding landscape of information.

When SonoSim initially came to us, they signed up to our self-boarding option. Once they tried this out, they saw many of the key features they needed were on the next tier of service and this was important to them, especially because in the health space, they need certain security and compliance requirements.

Our team engaged with them from there and the whole process took about three weeks. It was very, very smooth because they had already tried out the platform and built a prototype with it. Our engagement deepened as we provided tailored advice to ensure the integration could scale effectively with their business needs.

Phase 2: Implementing Vectara's Platform and Calibrating to SonoSim's Environment

The Vectara platform and process

AMR: Vectara's platform is at the forefront of generative AI and neural network advancements, providing exceptional search accuracy through semantic understanding. Our system excels in deciphering user queries, regardless of phrasing, ensuring the delivery of highly pertinent results.

For comparison, OpenAI and other models facilitate the fine-tuning of large language models. This process involves uploading proprietary data to iteratively

train a model, thus customizing its responses. Although user-friendly and efficacious, this approach is not without its challenges.

Fine-tuning begins with a base model, such as GPT-3 or GPT-4, already trained on publicly available datasets and various corporate contributions. The objective for businesses is to embed their unique dataset into the model. By doing so, the model evolves, integrating the new data to form a derivative, fine-tuned version, now enriched with both public and proprietary knowledge.

In this process, companies strive to exclude unwanted data through a technique called negative training. This involves instructing the LLM to disregard certain information. However, completely eliminating unwanted data is difficult. There might always be some residual data left behind, especially if there isn't a perfect understanding of what to exclude. To address this challenge, the new, customer-specific data is incorporated, not in its original form but as condensed, meaningful interpretations learned by the model. This way, the LLM leverages valuable insights while minimizing the influence of unwanted elements.

Fine-tuning's main drawbacks are its cost and computational demands. It's an intensive process that requires powerful hardware, like the NVIDIA A100 Tensor Core GPU, designed to provide acceleration for AI and analytics. It can take extensive time for the data to be assimilated, making it impractical for rapidly changing datasets.

Another concern is the potential for hallucinations. Even when fine-tuned, these language models lack strict mechanisms to prevent such errors, leading to occasional unreliable responses. In sectors where accuracy is critical, such as healthcare or finance, the inability to trace the model's reasoning process or cite data sources is a significant drawback.

Furthermore, even with careful training, bias and the inadvertent inclusion of copyrighted material remain risks, opening the door to, for example, a copyright owner coming after me instead of going after OpenAI, because I'm in a well-known brand using it. Intellectual property violations are a big concern for companies leveraging these models.

LLMS AND RAG ENABLE HYPER-PERSONALIZED EDUCATION FOR HEALTHCARE TECHNICIANS

In contrast, Vectara uses Retrieval Augmented Generation (RAG) to mitigate these issues. This method reduces costs, enhances security, lowers the risk of hallucinations, and allows for the provision of citations. It efficiently extracts, encodes, and indexes content from diverse formats (including PDFs, PowerPoint, images, video, and audio sources), aligning search results closely with user inquiries. It also enables the generation of contextually relevant and precise content.

Grounded Generation

How it works

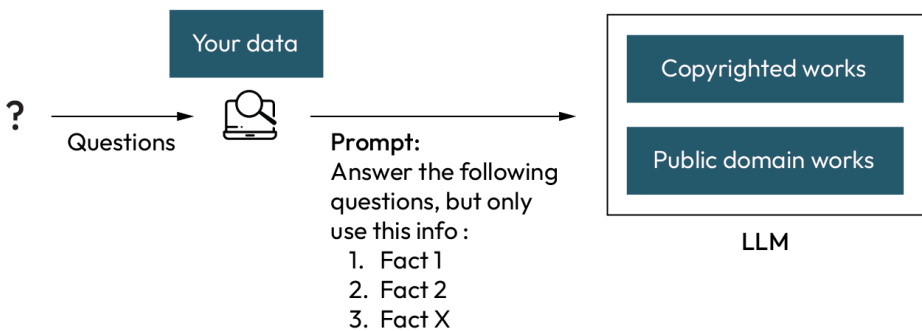


Figure 26.1 – Flow diagram demonstrating how grounded generation works

Our neural search technology, integrating hybrid search (keyword and LLM), significantly outperforms traditional keyword-based searches in delivering relevant results. This is because the engine operates not just on the recognition of keywords, but also on the comprehension of the meaning behind the words. The integration of hybrid search enhances the precision and context-awareness of the results. This multilingual capability also means it has a broad geographical reach.

RAG is an approach that leverages the power of LLMs to generate responses. The key aspect of this approach is that it retrieves relevant information from external data sources in addition to the user's data. This combined knowledge base allows RAG to tailor responses that are essentially a reflection of the user's

LLMS AND RAG ENABLE HYPER-PERSONALIZED EDUCATION FOR HEALTHCARE TECHNICIANS

unique knowledge base, making the interaction more relevant. That's really the secret sauce here.

This is all made possible by advanced information retrieval systems and machine learning models that prioritize semantic correlation over keyword matching. This semantic approach mirrors the cognitive processing of LLMs, utilizing vector embeddings to represent meaning.




Figure 26.2 – Vectara's end-to-end platform for retrieval, reranking, encoding, and generation

This novel technique is grounded in research, including papers authored by my co-founder, Amin Ahmad, which describe the development of neural networks that transition from linguistic to semantic comprehension. Despite currently utilizing GPT models as a foundation, we anticipate transitioning to Meta AI's LLaMA, a new suite of LLMs unveiled in 2023.

Translating language into meaning

Vectara's platform revolutionizes data processing by translating it into a vector database, essentially a "meaning database." This preserves the high-fidelity content, converting it into vectors that encapsulate the essence behind words, in whatever language you use. They are stored in their original content – not compressed or fine-tuned. So, when a query comes in, the engine matches the query's semantic intent with its knowledge base, ensuring only relevant, high-quality facts are used to construct responses.

Eliminating bias and enhancing accuracy

Our RAG approach prevents biases and copyright infringements by strictly using verified facts. It eliminates hallucinations, as responses are based on concrete data, not conjecture. This method is supported by advanced meaning engines, crucial for accurate information retrieval and ranking, thus empowering the LLM to generate precise, accurate, and reliable answers.

Responsiveness, balance, and copyright adherence

The process starts when a query is input into our search box. The system then collates results and synthesizes them into a coherent response. Crucially, the LLM is restricted to using only the provided facts, which are pre-validated and scored, based on whether the meaning matches that of the prompts, with 100% indicating a perfect match and 0% indicating no relevance. The model is instructed to cite these facts, employing a citation notation that reflects their relevance score. If none of the facts are pertinent, the model is programmed to refrain from fabricating information and instead admit to the lack of relevant data.

This method ensures a multi-perspective approach to information processing. When dealing with queries that inherently contain multiple viewpoints or possible answers, the system is designed to present a balanced view, incorporating all relevant perspectives rather than favoring any single one.

In addition, the model is constrained to only reproduce facts, which means it will not inadvertently disclose any copyrighted material unless it forms part of the facts it has been supplied. While the model may learn and replicate certain writing styles, it is our responsibility to ensure that these do not infringe on copyright laws.

Citations for transparency

Last but not least, citations play a critical role in the transparency of our process. Citations linked to the facts ensure explainability, allowing users to trace the origin of the information provided by the LLM. This significantly reduces hallucinations as the model consistently grounds its responses in factual data. Google and Microsoft also have ground generation.

The meaning engine

The key component to RAG's effectiveness is the capability of the meaning engine. If you don't find the right facts, the model collapses. It's wholly dependent on the facts being correctly ordered, based on their meaning and being scored correctly. This precision enables the LLM to perform adeptly even when presented with data it has not been previously trained on because of the strength of our zero-shot model learning capacity. Because it indexes new information, discerning the best meanings and making informed associations, it can deal with data categories or classes it's never encountered before. This advanced capability is foundational to our LLM's ability.

Specialization through verticalization

As we broaden our LLMs' verticalization, we tailor them to sectors like legal, medical, and accounting, while maintaining a general model for wider topics. Real-time prompt analysis helps us understand the customer's intent and guide queries to the most suitable specialized model. This focused strategy guarantees not only accurate but also contextually relevant responses, leveraging the specific knowledge of the selected vertical.

Simplifying integration

If you were building this yourself, you would have to worry about all of these steps. Building such a system from scratch involves numerous complex steps that can be daunting for the average engineer, let alone smaller firms without extensive tech resources. Vectara streamlines this with a simple API for data ingestion and interaction. We aim to reduce the heavy lifting, allowing users to easily upload their data and obtain responsive outputs without worrying about the back-end complexities. While we envision a vast connection framework, our immediate solution is a straightforward API that customers can use independently.

However, even if our API is very straightforward to get up and running with, there are still hoops to jump through, so that does add some friction. Comparable solutions are available from major cloud providers, such as Azure Semantic Search and Google Cloud's Vertex AI Conversation, while Amazon is expected to introduce a similar service, but we prioritized making our systems plug-and-play, requiring little to no configuration from users to achieve the best results.

Our focus on features such as security compliance, extensibility of answers and mitigation of hallucinations, bias, and copyright infringements, sets us apart in the marketplace.

The SonoSim Implementation

AMR: This was driven by the product team, with a relatively small engineering team. A customer support engineer was tasked to implement our technology into SonoSim's enterprise technology stack. This engineer, who oversaw the content management system across Sonosim's product line, was well-placed to support the implementation. Deep machine learning expertise was not a necessity; rather, the role required a developer who was conversant with SonoSim's systems.

We implemented our Retrieval Augmented Generation (RAG) based generative AI solution on their vast content repository of A-grade, optimally authored and trusted content. The solution allowed SonoSim's users to access content via a conversational interface. Instead of delivering a long list of results like a legacy search solution, a simple search began to yield highly relevant material instantly.

LLMS AND RAG ENABLE HYPER-PERSONALIZED EDUCATION FOR HEALTHCARE TECHNICIANS

Users could then effortlessly generate summaries, which, in turn, can be leveraged to create customized courses on the fly.

So, SonoSim could use our system for two things: the first, for search and information retrieval and, the second, for creating more helpful content for the technicians, which delivered a number of additional benefits, including:

- Custom course generation: Practitioners were able to generate five-minute courses with 600 words with almost no effort but a prompt. All information was drawn from our data and training repository, bypassing the need for any software development or course creation materials.
- Educator-facing features: These were also a significant part of the solution, as educators could design and roll out mastery tests to identify topics that practitioners found challenging or required further development. They could dynamically create training content to help individuals bridge their gaps and pass requisite skill and comprehension assessments. And administrators could track and monitor a user's progress over time.
- Tailored education: SonoSim has been able to hyper-personalize the entire learning experience for their users, something that was previously financially and logistically unfeasible. Generative AI and RAG played critical roles in helping break down those educational barriers.

Results and Impact

AMR: SonoSim's success hinged on a deep understanding of their user base – technicians who need real-time answers while operating ultrasound machines. Their whole business is about, "How do technicians find the right answers for what they're trying to do live?" So that was a no-brainer. This is the revenue. This user focus translates directly into their business model. They don't just sell machines; they offer a solution for in-the-moment education. They will tell you what's key isn't just the technology, but the extensive data library SonoSim has built. This library includes manuals, expert advice on machine optimization, and other crucial documentation. As they say, "The heart is the data, data that we have, which is all the manuals, all of the advice from experts on how to tune

the machines,” all of that kind of stuff they have in the documentation. That’s the power. And that’s true; getting these datasets can be very, very difficult.

Impacts included:

- **Efficiency:** The natural-language-based capability allows technicians to quickly find specific answers to their questions, reducing time spent searching through manuals and guides.
- **Learning and training:** The ability to query documents in a natural language format enhances technicians’ learning process. It’s a more interactive and engaging way to understand and learn how to use and calibrate various ultrasound machines.
- **Customer satisfaction:** Faster, smoother, and more accurate responses to queries mean improved user satisfaction.
- **Data utilization:** By leveraging generative AI technologies using their own data, SonoSim can make better use of their existing assets. The data that was previously underutilized can now be effectively used to train RAG implementation, leading to improved performance.
- **Innovation:** Implementing such advanced AI technologies positions SonoSim as an innovator in the field. It shows their commitment to leveraging the latest technologies to provide the best possible service to their users.

Key Learnings

- **Evolution of interfaces:** The evolution of software has gone through several phases, transitioning from command-line interfaces to graphical user interfaces, such as Windows 95, and subsequently to touch interfaces like those on the iPhone. Each phase brought software to a broader audience but necessitated the rewriting of applications to suit new modalities. Generative AI has now enabled AI systems to comprehend human language, marking a pivotal advancement in the interface between humans and software.

- **Data compression:** Similar to how the human brain compresses information, LLMs are trained on vast datasets while requiring only a fraction of the original data for storage. This data compression is valuable, but for company-specific data, referencing the original source is crucial. The model (X) processes the information, while another entity (Y) holds the actual facts.
- **Hallucinations:** AI hallucinations refer to inaccurate or misleading results generated by AI models. These errors can arise due to various factors, including limited training data, faulty assumptions made by the model, or biases present in the training data.
- **Fine-tuning:** The current LLMs available on the market, like GPT, are pre-trained on a diverse range of internet text and can be used by many users at the same time. However, they are not designed to learn or adapt to new information after training. This means that, while they can generate creative text based on the input provided, they can't be fine-tuned on a specific task without going through a re-training process. Fine-tuning builds on pre-trained LLMs by incorporating a company's unique data. This creates a customized model enriched with both public and proprietary knowledge.
- **Retrieval Augmented Generation (RAG):** RAG is a secure, cost-effective alternative to fine-tuning in LLMs. It reduces hallucination levels and uses semantic retrieval techniques for context provision. It can provide citations, enhancing the credibility of its outputs. It's adept at extracting, encoding, and indexing intelligence from various sources, which is then retrieved and ranked per user queries. It generates relevant content by considering unique contexts and historical materials. The choice between RAG and fine-tuning depends on the task requirements. Some projects may benefit from a combination of RAG and fine-tuning, enhancing model performance and reliability.
- **Knowledge retrieval:** LLMs significantly improve the efficiency and accuracy of knowledge retrieval within organizations, enabling swift and precise querying of proprietary information. LLMs can help bridge gaps in specialist education, especially where there are various learning modalities, high competency thresholds, limited time and resources, as well as teaching opportunities.

LLMS AND RAG ENABLE HYPER-PERSONALIZED EDUCATION
FOR HEALTHCARE TECHNICIANS

- **Trustworthy AI:** Trustworthy AI is a rapidly evolving field focused on developing reliable, transparent, and privacy-respecting AI systems. It has become an integral part of the AI system lifecycle to ensure beneficial and reliable outcomes. It requires governance and regulatory compliance, from ideation to design, development, deployment, and machine learning operations. In healthcare specifically, trustworthy AI is crucial for maintaining the integrity of medical education and practice. It addresses ethical concerns like patient privacy, the explainability of AI decisions, and minimizing bias.
- **Future-proof education:** As medical knowledge continues to rapidly expand, integrating AI ensures that educational content remains current and relevant. The ability to update and disseminate new findings quickly is key to future-proofing the education of medical professionals.



27 CASE STUDY

AI Personalization Increases Engagement in Nascent Tech Communities

IN CONVERSATION WITH HARSH SINGHAL: Data scientist and current head of AI and machine learning at Koo, a microblogging social media platform targeting non-English-speaking audiences with English-to-regional-language keyboards, local language newsfeeds, and hyper-local hashtags. Koo's AI-based "Topics" feature on the site allows users to select their areas of interest, enhancing content relevance, user engagement, and positive feedback.

LinkedIn: [harshsinghal](#)

Case Study in Summary

Case Study Classification

- Segment: B2C
- Industry: Consumer social media
- Geography: India, LATAM, Europe
- Value creation levers: Product, personalization
- Time to impact: 5 days-6 weeks
- Regulatory considerations: Medium

The Business Challenge

The social media landscape is dominated by a few global players, often leaving non-English-speaking communities underserved. Koo, a platform similar to Twitter, has a mission to cater to the 75% of the world's population who do not speak English by allowing them to share content in a variety of languages. With support for 20 languages, Koo focuses on enhancing user engagement through features like Topics, which is the subject of this case study, and personalized content based on user interests.

Koo boasts a diverse user base that extends beyond India to countries such as Brazil, Portugal, and Nigeria, reflecting its ambition to serve a global audience. The platform might be the first interaction these audiences have with social media, so Koo needs to meet them where they are in terms of tech literacy with a flexible approach.

Headquartered in India, Koo has expanded from its initial focus, supporting 10 Indian languages, to add an additional 10 international languages. Growth since its inception has seen over 60 million downloads across more than 100 countries.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

AI's inherent democratization of technology tools has enabled this Indian startup to take on Twitter in a cost-efficient way.

To drive engagement, Koo needed a scalable mechanism to personalize content to its multilingual users, addressing the existing bias toward English-speaking digital natives on global platforms.

The challenge included:

- Develop and launch a feature to encourage users in predominately nascent digital populations to become regular users on a social media platform and choose topics of interest to drive higher engagement metrics.
- Leverage AI to categorize multi-language content into relevant topics at scale.

The Response

Harsh built and launched “Topics,” an AI-based feature on the platform to enable users to choose their topic areas. This made the platform’s content more relevant, improving engagement and favorable customer feedback for the brand.

The work was done in three phases:

Phase 1: Determining prioritization and design. This phase prioritized the development of a new feature called “Topics,” to help users find content specific to their interest areas. To maximize efficiency, it was decided to initially launch in Hindi and English, the languages with the highest volume of users. To ensure quality, the team trained LLMs on human-labeled data, enabling the categorization of content into various Topics in other languages by leveraging research of the Indian Institute of Technology (IIT) Madras.

Phase 2: Building proof of concept. This involved rapid testing and iteration. A prototype was developed and tested among 40 internal users. Once the feature worked 80% of the time, it was taken to market with a rapid testing and iteration process. Community leaders and engagement teams conducted multilingual testing of the Topics feature. Human observation played an important role,

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

with the output of any model or algorithm being reviewed before being taken to production, in addition to standard model evaluation metrics. Post-launch, user feedback was actively sourced, providing valuable insights for further iterations. Automated testing tools were used to increase efficiency and accuracy in the testing process.

Phase 3: Deploying models and safeguards. This phase used open source software from Uber called Ludwig for rapid model training. BentoML was used for MLOps and model deployments. To ensure cost-effectiveness, there was tight coordination with the DevOps and back-end teams, with mindful usage of GPUs and CPUs, as well as use of real-time comparison and batch processing. The “Topics” feature was initially launched in English and Hindi, then expanded to other Indian languages. Safeguards were put in place to ensure content security and moderation through legal compliance, protection of personal data, robust content moderation measures, and anti-toxicity models to remove hate speech.

Case Study in Detail

Harsh Singhal profile

Harsh Singhal is an accomplished data science leader with over 15 years of technology experience, including leading teams of data scientists and other technical professionals at LinkedIn and Netflix. Harsh joined the Koo leadership team in late 2021, to build out Koo’s machine learning and AI practice as well as lead all data science and machine learning teams. His specialist areas are natural language processing and computer vision.

The Business Challenge

HARSH: Koo’s goal is to create world-class social media experiences for its target audience – individuals who do not consume content in English.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

To realize this goal, I was tasked with developing innovative user engagement features to enhance the personalized experience for our diverse user base not just in India but also in Brazil, Portugal, Nigeria, and others.

Knowing that other countries, such as China, could build large social platforms, and make them relevant to global users, inspired me to join Koo's mission.

Context-dependent issues

Addressing existing global platforms' bias toward English-speaking digital natives

HARSH: We can often take a digitized life for granted, but there are millions of people around the world who are just beginning their online lives. For instance, hundreds of millions of people who speak only Indian languages like Hindi, Tamil, Telugu, Kannada may have access to WhatsApp and Facebook, but are not widely exposed to micro-blogging or social media where they can express themselves in their own language.

The handful of global players that dominate the social media landscape leave non-English speaking communities underserved. So, providing a microblogging platform for users to express themselves in their native language is at the core of this business challenge, but for long-term success, you need content that aligns with their interests to keep users coming back.

Chicken-and-egg problem of getting the right VIP creators

HARSH: In the early stages of launching a social media platform, if there aren't enough users interested in a particular topic, such as sports, then even if a sports celebrity joins, they might not find an audience. This lack of engagement could discourage them from creating more content. Therefore, it's crucial to have a strategy to identify and onboard content creators who are relevant to the users' interests. The goal here is to create a self-sustaining model where personalized content attracts users, which in turn encourages VIPs to join and create relevant content. This engagement then attracts more users and creators in a virtuous cycle.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

Approach

A three-phase approach was taken.

Phase 1: Determining prioritization and design

Feature development and language prioritization

HARSH: My team set out to build a new feature, which we named Topics, that would allow users to explicitly state their interests, for example, entertainment, sports, or politics, when onboarding. We needed a way to map users to interesting profiles and determine what types of celebrities, influencers, and thought leaders to onboard, which in turn would generate and present relevant content to our users.

When faced with a project like this, the key first step when focusing on speed and efficiency is to reduce the complexity. So, one of the difficult decisions I had to make at the start was to choose between taking the feature (Topics) live in all languages versus the top one or two by volume of users, which were Hindi and English. While we may have lost momentum on other languages, the decision to prioritize these two was the right one because if we'd done it all together, we would have delayed the entire launch. Building it in one or two priority languages enabled us to scale it out afterward.

Data labeling and quality assurance

HARSH: While English and Hindi were prioritized, the team trained LLMs on human-labeled data to categorize Koo content into topics across 10 Indian languages and 10 international languages.

To get the data to the right level of readiness, we deployed an in-house team to review and manually label content in a process called human-in-the-loop (HITL). Our in-house team collaborated with our community leaders (who build and manage our language communities). Over several weeks, one leader from each language category reviewed and manually labeled the content.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

We developed Topic classification using Natural Language Understanding (NLU) technologies. These technologies include LLMS, for example, BERT, and more recently fine-tuned open source LLMs like Meta's Llama2.

This process produced a few thousand pieces of labeled content per category, which was our "gold standard" dataset. This was used to train the model, which was then used to classify the content on the platform. Often, in this process, the models can't classify some content with certainty, so it's then taken back to the original auditor to provide a clear label. This quality assurance helps the model improve over time.

Open source collaboration for model development with NLP

HARSH: To enable Topics for Indian languages where Natural Language Processing (NLP) research was nascent, we collaborated with top Indian research groups like the Indian Institute of Technology (IIT), Chennai, and AI4Bharat to come up with a solution.

I am very proud to say that our language model (nicknamed KooBERT) was open sourced on Hugging Face to spur further research.

Phase 2: Building proof of concept

Rapid testing and iteration

HARSH: Our culture is similar to the early stages of Meta, where we move quickly but with a heightened sense of awareness. Timeline constraints are coupled with limited budgets, so for proof of concept, I recommend getting 20 people in your company who use the product, and then dogfooding (testing the technology on ourselves internally for quality and experience before a wider audience) the results of the model in Excel or in a simple spreadsheet as a small proof of concept to see what it looks like. This means giving them the output that they would see on the platform if they were users. Then, get their feedback. If you do it for 20 people and it works 80% of the time, I think you have something there. Take it to market with a rapid testing and iteration process.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

We went through a wide range of tests, including following new Topics, how the Topic presents in settings, does the content under the Topic look right, plus the number of miscategorized content instances.

Our ML team and QA team did this for well over a week, around the clock. Given the number of languages involved, we didn't have people speaking all the languages on the team, so we roped in friends and families to do walk-throughs and test it in different languages. This involved a fair bit of coordination but there was a lot of positive energy around this as it was an exciting, high-impact feature for Koo. Some things we learned included:

- **Role of human observation:** Unless we looked at the output of any model or algorithm, we didn't ship it; we didn't take it to production. Of course, standard model evaluation metrics, (accuracy, false positives, ROC curves, and so on) are fundamental. But once past that, we look at the content that the model is predicting or classifying or recommending to see if it looks right.
- **Creating a set of internal power users:** Every time we had a new algorithm or feature release, we checked the feature among employees, across our teams (Data Ops, QA, and ML) who were all engaged users of the platform. They test a feature pre-launch as a user would experience it. This gave us a powerful pre-launch feedback loop.
- **Sourcing the voice of the customer:** After a few weeks post-launch, we actively sourced feedback from our users, which provided further, valuable color on how we could iterate and make the product better. While we are always very metrics-driven (with the model, engagement, KPIs, product performance, and so on), this human element of quality assurance tells us stories; for example, we learned that there was a strong sub-group that was pleased to find poets on the platform, so this became an interest area we introduced on Koo. We also used automated testing tools to increase efficiency and accuracy in the testing process.

AI PERSONALIZATION INCREASES ENGAGEMENT
IN NASCENT TECH COMMUNITIES




Figure 27.1 – Snapshot of KOO Topics features (Source: kooapp.com)

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

Phase 3: Deploying models and safeguards

Using open source to help launch multiple models at scale

HARSH: One of the key considerations with this business challenge was one of scale. In the beginning, it wasn't clear how we could deploy LLMs in production so we could serve millions of users at low latency levels. I think the science was there, we knew how to do it, but we wanted to do it rapidly and push six to ten models into production for the launch of a single language. And, as is often the case with high-growth companies, we didn't have a very large team (at that point, there were three or four people on the machine learning team, albeit with a mandate to grow). To solve this problem, we used open source software from Uber called Ludwig for rapid model training. Additionally, BentoML was used for MLOps and model deployments, with the teams working closely to integrate contributions and innovations for rapid deployment.

I'm very proud that a lot of the work we did with the open source community made great leaps with the Indian language NLP. Working with BERT, a foundational technology used in building LLMs that Google came out with in 2018 before GPT improved search performance enormously. Using BERT-based language models, we open sourced our BERT language model that was trained on Indian language and it's now available for anyone to go and use.

We also continue to have regular brainstorms and workshops with the teams so they can share their ideas and insights on how to continue to leverage open source technology effectively.

Continuing to contribute to the open source community not only enhances the company's reputation at the forefront of technological advancements but also helps foster innovation.

Coordinating with DevOps and back-end teams

HARSH: Tight collaboration between DevOps and back-end teams was also crucial to get our resources in order and ensure we were being frugal in terms of our cloud costs. We had to be mindful about how we used GPUs and CPUs,

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

when to run in real time versus batch processing (the latter is more economical in terms of costs), how do we use CPUs? When do we actually run it in real time? When can we do batch processing and save on costs?

Launch and initial scale

HARSH: So, we initially took Topics live in English, because a lot of the machine learning, NLP technology for extracting keywords for classifying content in various categories for training models is very mature for English. In Hindi, we had to struggle a little bit and that's where we innovated.

We worked closely with the team at AI4Baharat, India's most sophisticated NLP research group, supported by Nandan Nilekani, who is almost like India's CTO. We leveraged a lot of the work the team had published in research journals to develop our Topics in not just Hindi but eight more Indian languages, including Tamil, Telugu, Kannada, Marathi, Assamese. We took it live in our product, which gave us a lot of great engagement with time spent on the app increasing as a metric that we tracked.

Safeguards

HARSH: Ensuring content security and moderation is something we're passionate about at Koo. Everyone involved in the business wants to build Koo as a safe haven for our family members to use, so we've been very thorough with our safeguards creating the following models:

- **Legal compliance:** We are fully compliant with the laws of every country we operate in, such as the IT Act in India, which is akin to abiding by banking laws.
- **Protection of personal data:** Our Personally Identifiable Information (PII) is encrypted and accessed strictly on a need-to-know basis, subject to rigorous checks and balances, and usually involves our founders for particularly sensitive cases.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

- **Content moderation:** We've implemented robust measures to instantly remove NSFW content. Nudity and violence are not permitted in videos, profile pictures, or any uploads. Our system is designed to detect and delete such content within milliseconds.
- **Toxicity and hate speech regulation:** We've developed anti-toxicity models to identify and remove hate speech, based on legal definitions rather than our discretion. We use both automated systems for clear-cut cases and human oversight for more nuanced situations.
- **Transparency:** We maintain a transparency center to report the number of takedowns, whether due to government orders or violations of our platform rules.

The safety of Koo is a mission deeply personal to me, especially as I'm heavily involved in developing the machine learning that powers these systems.

Thoughts on hiring and developing talent

HARSH: In addition to all the product-related learnings that came out of the development, build, and launch of Topics for Koo, I would like to share some specific insights on hiring and developing talent for startups.

Who to target

Given the high demand for the data science skillset, like many leaders, I have faced challenges in hiring suitably experienced candidates because it's not possible to compete with much larger, well-funded companies when it comes to compensation. With this in mind, we often look for relatively junior but motivated technical developers to join the team, focusing on finding passionate self-starters who are aligned with Koo's mission.

How to find them

I developed an interview process focused on practical assessments like reviewing/building models hands-on, to gauge immediate technical skills. I frequently ask candidates to submit ideas for features they would build if hired to filter for those invested in Koo. In my view, this is a particularly strong indicator because

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

someone who is motivated and has spent time deeply exploring the platform and coming up with ideas for high-impact features is highly likely to emerge as a strong future contributor to the team.

Developing and retaining talent

Hiring is just the first step. To develop and retain talent, I pair junior hires with senior team members as mentors to help them upskill and transition from amateur coders to professional developers. I have also instituted structured and deliberate processes like writing design docs and retrospectives to promote continuous learning.

Results and Impact

HARSH: After rigorous assessment and QA, the Topics feature was launched in Hindi and English first, delivering a minimum viable product to gather initial feedback. The accelerated launch increased time spent on Koo by 100% for users that adopted it. With Topics established, the team could then scale the feature across all 20 languages.

Topics was key for Koo to build thriving, localized communities and personalize the experience. Understanding the popularity of Topics enabled Koo to onboard relevant celebrities, thought leaders, and creators that users wanted to follow. Advanced features like personalized recommendations based on user affinities further connected users with relevant voices to follow.

The visibility Topics provided into audience interests also helped creators better reach a more targeted audience and encouraged them to create more engaging content. Increased time spent showed Topics helped convert first-time users into an active user base.

In summary, our approach is a cyclical process: an increase in creators attracts more users, which boosts on-platform engagement. This leads to a growth in followers, reaching a critical mass that allows for the development of more

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

categories. These new categories, in turn, attract even more creators, perpetuating the cycle. This self-reinforcing mechanism is our “flywheel.”




Figure 27.2 – Koo flywheel showing the virtuous cycle of increasing engagement (Source: Koo)

The impact of this iterative approach resulted in doubling the time spent on the platform by users.

Key Learnings

- **AI-powered language support:** Developing an AI-powered multilingual system promotes much broader, inclusive, and diverse adoption. By implementing automatic translation features and collaborating with native speakers or language experts, platforms can overcome language barriers to access information.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

- **Human-in-the-loop (HITL):** HITL is the combination of human oversight with AI in an iterative workflow to build robust and accurate AI systems. This process strikes a balance between automation and human expertise for machine learning. For instance, when building a corpus of content to train a model, human expertise can ensure accuracy through quality assurance processes.
- **Dogfooding:** Internal testing, otherwise known as “dogfooding,” is the process of testing one’s own products and services for quality and user experience, before releasing them to the public. This process leverages internal “super users” and has a feedback loop to developers that allows them to identify bugs and usability issues early in the development cycle.
- **Role of DataOps:** The DataOps team streamlines the process of building, testing, deploying, and managing data products. Its purpose is to implement rigorous testing procedures for quality assurance to ensure high-quality data throughout the AI lifecycle.
- **Recruit the talent:** Hiring technically skilled, self-motivated individuals passionate about problem-solving ensures a committed team. A recruitment process should evaluate technical expertise, alignment with company values, and intrinsic motivation. Highlighting learning and development opportunities can help attract and retain talent.
- **Open source:** Leveraging open source technology offers fast-track innovation and rapid deployment. Contributing to the open source community also builds a collaborative environment and reputation.
- **Optimize resource allocation:** Making smart resource usage decisions, such as determining when to use GPUs and CPUs, and determining whether to execute processes in real time or perform batch processing, can significantly reduce costs associated with AI development.

AI PERSONALIZATION INCREASES ENGAGEMENT IN NASCENT TECH COMMUNITIES

- **Customer feedback:** Gathering feedback from customers post-launch through surveys and customer support queries provides valuable insights into customer needs and preferences. Using a metrics-driven approach to track model performance, engagement, KPIs, and product performance helps to measure and improve the product.



28 CASE STUDY

AI-Powered Virtual Agent Augments Service Efficiency

IN CONVERSATION WITH PAT CALHOUN: CEO and Founder of Espressive – a software company that is transforming enterprise self-help with AI. Espressive uses AI technology to deliver enterprise-ready digital workplace assistance that reduces cost while boosting employee experience, productivity, and satisfaction.

LinkedIn: [pacalhoun](#)

Case Study in Summary

Case Study Classification

- Segment: B2B
- Industry: Technology
- Geography: Europe, MENA, APAC, AMER
- Value creation levers: Operations
- Time to impact: Less than 6 months
- Regulatory considerations: Low

The Business Challenge

Espressive's client, DXC Technology, is a fast-scaling global technology services and insurance software provider. DXC was experiencing a high volume of calls to their IT service desk, resulting in long queues and a poor employee experience, while also keeping the internal IT team away from strategic initiatives. The challenges included:

- Increased service desk dependence with a sharp rise in, and continued prevalence of, remote working.
- Service desk agents inundated with repetitive calls.
- Disparate IT systems meant a solution needed to integrate with ServiceNow, SharePoint, and other data sources.
- Integrated solution must auto-escalate unresolved issues to human teams.

The Response

DXC Technology selected Espressive, which develops AI-based digital workplace assistance for enterprise businesses, to implement a virtual agent system alongside a system overhaul and to provide 24/7 instant resolution to employees' questions. The client sought a way to create an automated solution with a filtered option toward live help if needed.

The work was done in three phases:

Phase 1: Scoping the requirements. The initial goal, to decrease repetitive questions to the service desk, was achieved in this phase by understanding what employees were asking in their day-to-day language and analyzing and automating the organization's top call drivers.

Phase 2: Developing the solution. This phase developed the Espressive virtual agent, which DXC called EVA (Everyone's Virtual Assistant), and integrated it with a proprietary centralized workplace solution, to provide employees with streamlined support.

Phase 3: Integration and adoption. This phase integrated EVA with multiple service desk back-end systems and other data sources. Adoption was prioritized with a dedicated global program to build employee confidence.

Case Study in Detail

Pat Calhoun profile

Pat Calhoun is a visionary leader with an intense focus on user experience and customer adoption. He has founded two companies: Airespace and Espressive. As CEO at Espressive, Pat is set to transform the enterprise self-service experience to a consumer-like approach that drives employee adoption and significantly reduces help desk calls. Pat's first startup, Airespace, grew revenues to over \$80M in two years before being sold to Cisco for \$450M. Most recently, Pat served as senior vice president of product at ServiceNow, where he was responsible for

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

ServiceNow applications. Before that, he was general manager of the McAfee network security business. Pat also served as CTO for the Cisco \$14B switching, routing, wireless, and security access business, and GM of the Cisco Identity Services business. Pat holds 35 patents and has been published in more than 16 publications.

Espressive, a pioneer in automating digital workplace assistance with generative AI, redefines how employees get help by delivering exceptional experiences across the enterprise. Espressive Barista, the company's innovative virtual agent platform, takes on the role of the service agent, bringing the best of human experience with the best of artificial intelligence. Barista leverages a proprietary and domain-specific large language model (LLM), as well as integration with generalized LLMs, to automate the resolution of employee questions, issues, and requests with personalized experiences that result in employee adoption of 80% to 85% and reduced help desk call volume by 50% to 70%.

The Business Challenge

DXC Technology, a global IT services and consulting company headquartered in the United States, helps global companies run their mission-critical systems and operations while modernizing IT, optimizing data architectures, and ensuring security and scalability across public, private, and hybrid clouds. The world's largest companies and public sector organizations, including over half of current Fortune 500 companies, trust DXC to deploy services to drive new levels of performance, competitiveness, and customer experience across their IT estates.

With over 130,000 employees spread over 70+ countries, DXC is a global company with employees requiring 24/7 IT support. Prior to their involvement with Espressive, DXC employees experiencing IT issues would call the service desk to speak to a human agent, who would either walk them through solving the problem or, if more help was needed, create a ticket for them directly in ServiceNow. Employees could also email the service desk, which again would lead to an agent creating a ticket. Without any automation, even routine questions like password resets required human interaction. The result was long queue lines even for simple questions, keeping service desk agents from focusing on more strategic initiatives.

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

Kristie Grinnell, Senior Vice President and Chief Information Officer at DXC, realized that manually handling every IT issue was not only inefficient for her team but also created a poor employee experience. To maintain support across multiple offices, countries, and time zones, they needed a centralized solution.

DXC's Experience Platform, DXC Uptime™, a proprietary centralized workplace solution that empowers employees to connect, collaborate, and work seamlessly, was rolled out. Along with Uptime, DXC Digital Support Services provides personalized support so employees can focus on their respective work, not troubleshooting IT issues. Uptime provides employees with a simple way to interact with technology and gives IT a consolidated view of systems for easier asset management and more effective service delivery.

However, consolidating the IT experience only solved part of the problem. Grinnell and her team recognized that they needed to bring in automation if they were going to decrease the number of calls going to the service desk. They needed a solution that provided a fast return on investment, so they decided to automate employee self-help with a virtual agent.

DXC chose Espressive to supply DXC Digital Support Services with a consumer-like, AI-based virtual agent to embed into their new centralized workplace solution, providing 24/7 instant resolution to employee questions rather than directing employees to navigate knowledge articles and find their own solutions. Espressive committed to a minimum deflection rate of 40%, providing DXC with confidence in the potential for significant operational efficiency improvements. Additionally, adherence to stringent global security protocols, such as GDPR, was a key selection criterion. Espressive effectively met DXC's rigorous security standards.

Context-dependent issues

DXC had an outdated way for employees to get IT help. Employees had to talk to a human agent every time, even for simple requests. As DXC continued to expand globally to over 130,000 employees, this resulted in long wait times, keeping service desk agents from focusing on more strategic initiatives. Automating top call drivers, via the virtual agent, helps keep these repetitive questions from the service desk.

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

The solution needed to integrate with ServiceNow, SharePoint, and other back-end data sources to automate top call drivers and provide internal knowledge.

If the virtual solution didn't know an answer to an employee question, it needed to automatically create a ticket within ServiceNow and route it to the correct team for human chat.

Approach

A three-phase approach was taken.

Phase 1: Scoping the requirements

PAT: Lots of businesses recognize the need for automated systems (chatbot or virtual agents) to reduce costs. Many have deployed self-help platforms, which are essentially knowledge search bars, but, as we've found, these often fail to provide valuable help as people don't like to read instructions (they prefer things to be fully automated), so employees will just pick up the phone and call the service desk.

ChatGPT has changed employee expectations around how virtual agents are supposed to work, leading to more complex and varied queries. However, what I repeatedly see is a lack of skills to deliver them because of a complete misunderstanding of the effort level that it takes to build a language model that really understands human language.

A recent study by an analyst firm found 25% of the cost is the build and 75% is the maintenance. And I can vouch that this is true. It's not just building the bot; it's also the continuous improvement needed on it.

So, our philosophy out of the gate was how can we build a virtual agent on a conversational platform that really understands the language of employees? We built the language model and then continued to improve it by focusing on that use case.

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

ChatGPT has also led to a desire for companies to expand to enterprise search; however, these technologies often lack the context of the enterprise and its tools, leading to incorrect answers. So, Espresso focuses on higher volume, low-hanging fruit when we come into an organization, doing an analysis of their top call drivers.

Focus on commonality

PAT: The virtual agent is agnostic because what we ended up discovering was that it doesn't matter who you work for; 85% of what employees ask for when calling the service desk is the same. Each business will have its unique lingo, acronyms, or buzzwords, but the things that they're looking for, the problems that they're experiencing, are universally the same.

For example, "I need a new laptop," "I need access to a piece of software," "What are the company vacations?," "I have a question about my retirement plan." It's all the same. So, focusing on a domain that was universally the same across all of these different companies provided us with an opportunity to really focus on improving our language model so that when we get deployed, we have a high degree of understanding. And then, of course, with additional oversight on supervised capabilities, we continue to improve that over time.

Phase 2: Deploying the solution

PAT: We built the Espresso virtual agent, which DXC called EVA (Everyone's Virtual Assistant), and rolled it out with Uptime™ so DXC employees wouldn't need to be trained twice on how to get help. Rather than calling or emailing the service desk, employees now go to Uptime™ to use EVA to get instant resolution to their questions, 24/7.

If EVA doesn't know an answer, the system automatically creates a ticket within ServiceNow and routes it to the correct team for chat – without human intervention. Deflection is a term used for tickets or issues that are resolved with no human touch. But for complex issues, there's a seamless handoff to a human.

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

Security

PAT: DXC has strict security standards, so a partner who maintains compliance and certification with the world's most stringent security protocols, such as adherence to GDPR, was critical. DXC required Espressive to take this a step further and keep all user profile information anonymous. I'm proud to say that the close working relationship between DXC and Espressive meant we successfully did this before launch.

Hans Meeldijk, the Senior Program Manager at DXC, said, "We see the Espressive team as a partner, and they treat us like one. They go the extra mile for us."

Phase 3: Integration and adoption

PAT: EVA integrates with systems like SSO, asset management systems, and ServiceNow, and is always increasing the number of enterprise platforms it can integrate with out of the box. For example, for another customer, we recently integrated it into the Ariba infrastructure, which allowed employees to interact with our virtual agent instead of having to become experts in Ariba.

Adoption

PAT: Employee adoption is always one of the biggest risks in this kind of transformation. DXC employees were used to going to a live agent for help, so we needed them to trust that their issues would be solved or routed correctly by a virtual agent instead of a human. To tackle this, we built employee confidence by ensuring EVA seamlessly answered top call drivers.

The DXC and Espressive Customer Success teams worked together to automate password reset, VPN issues, checking ticket status, software provisioning, and more. In addition to promoting the benefits of EVA, DXC utilized the Espressive Employee Adoption Program, creating content that helped DXC employees understand why using EVA would help them better than calling the service desk.

Challenges, costs, and personnel

PAT: More and more organizations are interested in fine-tuning language models like GPT-3 with their enterprise-specific data. However, I find that they often underestimate the complexity and costs involved. It requires access to relevant data, an understanding of the model, and the ability to continuously update it as new data becomes available. Finding and organizing the necessary data can be a significant challenge, as content changes daily, and is often scattered across different sources.

There's also a need for skilled personnel, a trend I think mirrors what we found seven years ago when organizations wanted to build their own bots. Now, the focus has shifted to building enterprise search solutions using language.

Results and Impact

“UPtime and our digital assistant are rolled out to about 90% of our employees across the globe and we’re deflecting about 56% of tickets. We hope to get up to 60%. By continually driving automation and increasing the deflection rate, we’re driving cost savings, and improving employee satisfaction and productivity.”

— Kristie Grinnell, senior vice president and chief information officer at DXC

Impacts included:

- Within two months of going live, up to 56% of tickets were fully automated, on target to 60%.
- The virtual agent, EVA, was able to achieve 81% ticket deflection within a month of going live.
- Since its launch, EVA handles over 23,000 employee interactions globally every month, reducing the number of phone calls to the DXC IT service desk by 36%.

Future plans

PAT: EVA will soon be the holistic source for answering employee questions across departments. DXC has already started to expand EVA to include support for HR and supply chain management, with more to come. They are also integrating EVA to Microsoft Teams to make EVA even more accessible to all employees.

Key Learnings

- **Complexity of human language:** Many businesses encounter failure in AI development primarily due to underestimating the complexity of human language comprehension. This complexity often leads to significant maintenance costs post-deployment, accounting for 75% of total costs.
- **Optimize remote working support:** The growth in remote work has increased service desk reliance, increasing the need for automation to cope with demand. Users often prefer automation over reading instructions, further underscoring its importance in enabling remote work. Aiming for automated ticket resolution without human intervention can significantly enhance efficiency.
- **Virtual agent model design:** The design and integration of a virtual agent can significantly impact its effectiveness and applicability. Designing the agent to be sector-agnostic, capable of accommodating over 85% of common queries, can foster high understanding and response accuracy, making it a versatile cross-industry tool. Furthermore, focusing on the uniformity of service desk queries can facilitate rapid solution deployment for improvements in service desk operations.
- **Transition from virtual agent to human:** Ensuring a seamless transition to a human agent for complex issues is crucial for maintaining user satisfaction and trust. This involves designing the AI system for easy handoff when the issue surpasses the AI's capabilities.

AI-POWERED VIRTUAL AGENT AUGMENTS SERVICE EFFICIENCY

- **Updates and integration with familiar interfaces:** Regular updates and integrations enhance the model's capabilities, keeping it relevant and valuable. This includes integrating with authentication systems, asset management, and ITSM systems (like ServiceNow) ensure that the virtual agent remains relevant and useful in a rapidly evolving tech landscape. This provides a more seamless user experience to access the information and support users through a familiar interface.
- **Balance value and complexity of fine-tuning models:** There's value in fine-tuning language models to meet specific enterprise needs, but it's crucial to be mindful of the complexity and costs, and the need for skilled personnel for model fine-tuning.



29 CASE STUDY

Generative AI Creates a Paradigm Shift in Innovation Processes

IN CONVERSATION WITH DAVID BOYLE: Author and expert with two decades of experience in using technology for audience strategy in creative industries. He has influenced major brands and co-wrote *PROMPT for Brands* with Richard Bowman. This case study explores how he used AI to revolutionize the innovation process at a global consumer goods company.

LinkedIn: [beglen](#)

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Case Study in Summary

Case Study Classification

- Segment: B2C
- Industry: Consumer goods
- Geography: Global with brands in more than 100 countries
- Value creation levers: Product innovation
- Time to impact: 1-6 months
- Regulatory considerations: Medium

The Business Challenge

Given the complex challenges associated with innovating across diverse markets, a global consumer goods company turned to generative AI to enhance and standardize its processes. The aim was to create a system that could navigate the intricacies of different markets while fostering innovation. They needed to overcome embedded approaches and low mastery of skills to work with Large Language Models (LLMs). The challenge was to:

- Improve the speed of product innovation through AI.
- Develop a more streamlined, efficient, and effective system across diverse geographies, overcoming legacy approaches.

The Response

The project developed a new AI-assisted innovation process drawn from numerous case studies by global innovation teams working on diverse products. These studies offered insights into traditional innovation methods worldwide, whether the teams were rebranding or launching new products. The project was enriched

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

by these real-world examples, providing practical context to the theoretical use of language models in innovation.

The work was done in six phases. Four out of six phases, each leveraging LLMs, were completed, with the client still working through the final two stages:

Phase 1: Skepticism and hesitation involved overcoming initial doubts about AI's efficacy and practicality, due to previous over-hyped technologies.

Phase 2: Initial exposure and surprise involved exposing AI to C-level executives to demonstrate its potential in a specific domain.

Phase 3: Experimentation and trial involved small-scale experimentation to understand AI's capabilities and how it fits into existing workflows.

Phase 4: Realizing practical applications involved recognition of AI's role in enhancing creativity and decision-making, to deliver tangible outcomes and benefits.

Phase 5: Adoption and integration involved the standard integration of AI into everyday processes, becoming an essential tool.

Phase 6: Evolution and learning involved continuous learning and adaptation as AI tools evolve.

The five-month project, involving numerous stakeholders, resulted in a flexible, comprehensive 300-page manual for product innovation, detailing case studies from different countries and the application of LLMs at every step of the process. The global initiative demonstrated that AI-assisted humans outperform those without AI in generating successful innovations.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Case Study in Detail

David Boyle profile

David Boyle is an author with over 20 years of professional experience using technology to develop audience intelligence and strategy capabilities in the creative industries. His work has influenced the culture and economics of some of the world's most prestigious entertainment and luxury brands, including EMI Music, HarperCollins Publishers, and the BBC.

DAVID: The work we've been doing with generative AI is about synthesizing existing information and expanding on it, as well as uncovering gaps in our thinking. It's about using generative AI to augment our capabilities and make us more efficient and effective in our roles. LLMs are ground-breaking in identifying patterns and can generate innovative solutions when users are trained in how to use them. By adding proprietary data, often derived from qualitative, ethnographic research, we can gain even deeper, unique insights than those based on the LLM's training data alone.

In terms of synthesizing existing information for innovation, we're using AI to process vast amounts of data that would be impossible for a human to handle alone. This includes market research, customer feedback, and internal performance metrics. The AI identifies patterns and trends we might miss, helping us make more informed decisions. Additionally, AI's predictive capabilities allow us to simulate different scenarios and predict likely outcomes, helping us identify potential issues before they become problems. This proactive approach is a huge advantage in today's fast-paced business environment.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

We're finding that product innovation is not unique in its process requirements. We're coming to the same conclusions in almost every area we look at, across multiple processes from team management and planning to songwriting. Every single area we examine shows that the combination of human creativity and generative AI is more potent than either humans or AI alone. Technology can enhance human creativity and analysis, but the human touch is crucial for imagination, making counterintuitive connections, and initially asking the right questions.

For example, we recently ran a training session for a pharmaceutical company. Their brief was, "Can we do a better job of coordinating, planning, marketing, and strategizing with multiple partners using AI? And can we avoid hiring more people so we can stay lean?" My answer is absolutely yes because AI, while a massive help, is just one piece of the puzzle. We're still very much in the driving seat. It takes human insight to truly understand pain points and develop meaningful solutions.

Moving forward, every single business process needs to incorporate AI at its heart. Otherwise, you risk being slower, more expensive, and offering lower quality, with employees working harder and being less happy. The innovation process is just one area where we've applied AI and decided that it wins. However, we know that's true in the vast majority of business processes.

Specifically, we've deployed AI to help supply chain teams predict product demand and ensure the right stock levels in the right places. We've also supported marketing teams using customer behavior data to deliver more relevant and engaging content.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

The Business Challenge

DAVID: A leading consumer goods company came to us to help them leverage AI around research. The request was framed pretty much around, “Can you write some prompts for us that will enable us to do product innovation effectively in an automated way?” but they didn’t say the word automation.

I wasn’t aligned to full automation. Language models can, to a certain extent, be automated or semi-automated, but the essence of their function, and their real power, is as a tool to augment human capabilities. They enable us to perform tasks more efficiently, more accurately, and with greater clarity. But they need strong supervision. The notion that they could entirely replace human input was something we strongly disagreed with from the outset.

So, our philosophy was very much rooted in being able to help the teams use language models in the innovation process. Yet, throughout the whole project, there was this tension. On one side, there were voices advocating for full automation, and on the other, our steadfast belief in helping people use these tools, not replacing them with these tools.

By the end of the project, I think they completely agreed, and they’ve publicly acknowledged this saying that language models radically improve how people do their jobs. They also publicly stated that while AI is vital, focusing solely on technology and overlooking investment in people is a significant error. To foster innovation and growth, companies need to invest equally in humans and AI.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Approach

DAVID: We deployed four out of six phases in the journey of AI adoption, with the client currently working through the last two:

Six phases of AI adoption

These phases represent the journey from initial skepticism to widespread integration of AI tools.

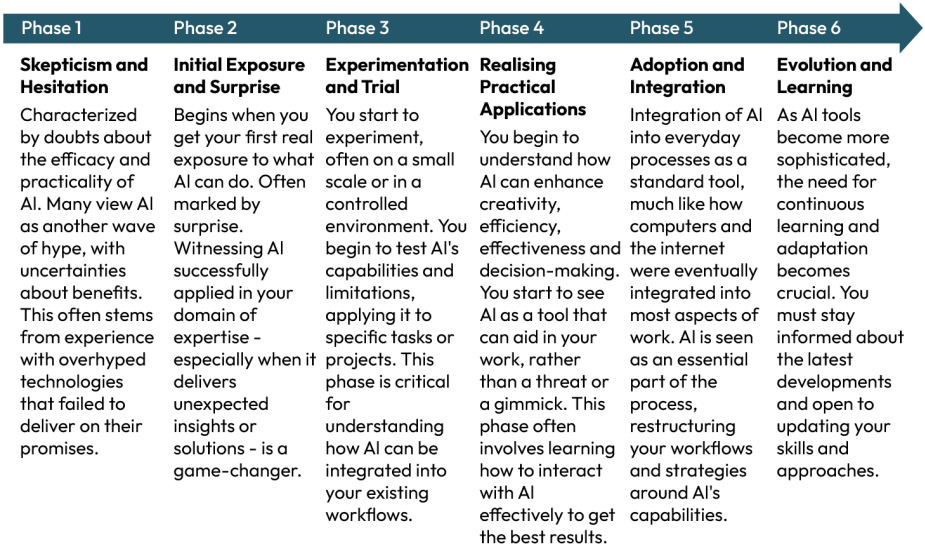


Figure 29.1 – Six phases of AI adoption (Source: Prompt.mba)

GENERATIVE AI CREATES A PARADIGM SHIFT
IN INNOVATION PROCESSES

As an organization progresses through the phases, they see greater and greater benefits from AI:

Phases of AI benefits

These phases represent the journey to benefits and transformation through AI tools




Figure 29.2 – Six phases of AI benefits. (Source: Prompt.mba)

Phase 1: Skepticism and Hesitation

DAVID: First off, skepticism and hesitation. It’s incredibly beneficial to maintain a healthy level of skepticism. If someone suggests a radical departure from a method that’s been in place for 50 years, it’s natural to be wary. It could easily be just a sales pitch, and it’s probably best to ignore it. There’s a lot of hype in this space, so skepticism is not only good, it’s necessary.

Phase 2: Initial Exposure and Surprise

DAVID: However, when you see a real example of one of your brands generated by AI right before your eyes, you start to think, “Okay, something’s happening here.” We demonstrated this early on in larger training sessions on broader themes, and they acknowledged, “Alright, this is real. Something’s going on here.”

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

So, when we came on board for innovation, they asked us to help them experiment with how it could work. This led to real-world use cases. We tried it with more brands and in more countries around the world. We learned more about the kind of problems it could solve and what good looked like in their eyes. We gathered more and more evidence, more and more belief that this approach was so different, so much quicker, better, clearer, and more fun than the old process, that the client decided to invest more.

We quickly achieved some truly remarkable results. The first notable result is when you use a language model to accomplish something that an expert deems to be exceptionally good. You can reach this point fast when working with a key decision-maker. That's what I love about working with CMOs or CEOs, when they say, "Wow, this is very impressive." You can then swiftly overcome that initial resistance.

But what we demonstrated was that in 45 minutes of careful, thoughtful work, you could come up with a new product innovation, which was really, really good. And there was no way to do that in the old world.

Phase 3: Experimentation and Trial

DAVID: This phase involved small-scale experimentation to understand the AI's capabilities and how it fits into existing workflows. It involved the evaluation of a multitude of real-world case studies involving actual innovation teams from across the globe. These teams were working on a diverse range of products, both food and non-food. These case studies provided valuable insights into the workings of teams in rural countries worldwide, who were still adhering to ingrained methods of innovation.

Sometimes, they were looking to enhance an existing brand, reposition it, and relaunch it in the market, or aiming to introduce new products. Regardless of their objectives, these case studies offered a glimpse into their processes, their challenges, and their unique approaches to innovation. This wealth of real-world examples served to enrich the project, providing tangible, practical context to our theoretical understanding of how language models can aid in the innovation process.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

The traditional innovation approach is a lengthy and comprehensive process. This becomes more complex in a global company. It involves consumer research, various stages of development, and then distilling that research through a series of workshops. This process takes on many forms and templates, with key milestones and meetings along the way. Essentially, it's a process devoid of AI, relying instead on market research, brainstorming innovation trends, and a great deal of creative writing based on these elements.

What we ended up doing was collaborating with real teams worldwide, who were accustomed to the old way, and assisting them in transitioning to the new, AI-supported way. In the process, we were defining what this new way entailed.

In our exploration, we conducted dozens of case study exercises, developed in partnership with the business. We discovered that out of a total of 150 prompts we developed, around 15 stood out as particularly useful. However, these prompts require personalization to the specific project, meaning you can't just apply them blindly – you still have to engage your intellect.

One thing we've done with the prompts is to be explicit about what we want and the format we want it in. The clearer you are in the prompt about your requirements, the clearer the responses. This approach can lead to immediate product concepts.

I believe it's incredibly crucial that the human working with AI gets to view the prompt. This is the challenge I have with custom-built interfaces – I can't view the prompts. If I read someone's custom-generated meeting summarization but can't view the prompt, I can't easily work out what it is doing (and not doing). Is it seeking insights or nuggets or catchy phrases or bold phrases or timid phrases or action phrases? Is it seeking specific actions or just insights? These are all aspects that I might care about or want to modify. I genuinely believe concealing prompts is an error.

We tested whether we could go straight to creating a solution or inventing a product in a single step and without feeding any background knowledge into the language model. The answer is that it's possible, but the results are far from

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

excellent, and a long way short of outstanding. The idea of automating the process is still elusive.

Jumping to a new product innovation in a single prompt does get you lots of thought starters, but this way it's challenging to pinpoint where the responses go wrong and how to fix them. Essentially, our conclusion is that for any reasonably sophisticated requirement, you need to follow some kind of multi-step process when using AI. You can't just leap to the answer. You don't have to go through the entire process, but you should certainly follow some kind of procedure rather than jump straight to an answer in a single prompt. That's a significant lesson.

Phase 4: Realizing Practical Applications

DAVID: Phase 4 involved enhancing the innovation process to realize tangible outcomes and benefits. This phase was designed to ensure that the integration of AI into our processes was not just theoretical, but also practical with measurable outcomes. The steps were meticulously planned and executed, keeping in mind the ultimate goal of enhancing the company's innovation capabilities and success rate.

In this approach, we concentrated on four key skill areas essential for leveraging LLMs. Over the past year, our training programs have been instrumental in helping individuals develop the foundational skills necessary for working effectively with any LLM. We believe these skills are exactly what people need to harness the full potential of LLM technology. The core of enhancing human capability lies in acquiring foundational skills. Once these skills are acquired, the particular tools employed – whether it's Otter.ai for summarization or Copilot for coding – become less important, as these tools are bound to change over time. However, the key is to understand their underlying operations. Developing the ability to craft effective prompts, modify them, iterate, and manage data inputs is essential. With these skills, you can adeptly navigate any large language model. Essentially, our focus is on imparting these fundamental competencies.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Worked examples of each LLM skill



Condensing information

1. **Basic:** summarize this article/report
2. **Good:** synthesize this diagram
3. **Better:** help me structure my thoughts
4. **Advanced:** create meeting minutes



Desk research

1. **Basic:** help me understand my audience
2. **Good:** generate deep audience insights
3. **Better:** provide facts for a new market
4. **Advanced:** consolidate market entry info



Content generation

1. **Basic:** communicate meeting minutes
2. **Good:** draft a follow-up email
3. **Better:** draft an email in another language
4. **Advanced:** create an internal plan



Creative thinking

1. **Basic:** come up with new product ideas
2. **Good:** position this idea for my audience
3. **Better:** help me build my plan/argument
4. **Advanced:** critique my plan/argument

Figure 29.3 – Critical key skill areas for LLM. (Source: Prompt.mba)

I believe that any team needs to master how to use LLMs across all four areas. Yes, people might use one or two, but they often don't realize the existence of the other two. So, these are four relatively comprehensive skills.

In this fourth phase, we followed the following specific steps.

Step a. Clarity of desired outcome

DAVID: The first step is being very clear about what it is you're actually trying to achieve asking, "What are my goals for this new product innovation? What does good look like? What does it have to satisfy in terms of business goals or consumer goals for this to be a success?" Applying the logic Einstein suggested about solving a problem, if I had an hour to solve a problem, I'd spend 55 minutes working out what the problem was and 5 minutes solving it. It is the same with language models. So, get to work on what problem you're trying to solve first of all, before even starting on solving the problem.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Step b. Understanding consumer needs

DAVID: The second step focuses on deeply understanding consumer needs effectively and then understanding what consumer segments have those needs. Then, what unmet needs may be in the category?

When starting to explore audience needs, if you have what we call prep materials to include, you can put them in here as part of your prompts. Processing input materials that you already possess with AI is still a win. You can process a significantly larger number of documents and synthesize them more rapidly for the context of the product and the category you're working on.

If you lack that consumer research, that's not a problem. You can extract useful consumer insights from the AI's training data, given that it has read a vast number of human reviews of products and services and possesses a deep understanding of underlying human needs. Therefore, there's another section of the client manual we produced that says, "Don't have any inputs? No worries, don't bother reading the chapter on processing inputs. Read this chapter instead." And it provides the types of prompts you can use to get the AI to provide you with useful human insights, audience segments, and so on. So, it's a very flexible process with numerous different approaches.

This brings us on to the second LLM skill area of desktop research. This is essentially leveraging the AI's training data and saying, "Hey, help me find insights, facts, unmet needs, or similar things." But again, this applies equally to pharmaceuticals, like when you say, "Help me understand the pharmaceutical and healthcare market in Israel." So, desktop research is a universal skill.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Step c. Innovating new products

DAVID: The third step is formulating a product off the back of the outcomes of the initial phases, being clear about what unmet needs it serves and the reasons why people should be excited about that product. Then, write out the full concept in consumer-facing language.

This step requires content creation skills, bringing the segments to life, asking the AI, “Can you narrate a story about this concept? Can you enhance this audience profile? Can you provide me with some scenarios to illustrate this concept?” Supplementing elements like these with good writing, style guidance, and inputting the pieces that you’ve developed from previous steps will create compelling concepts.

If you already have your product benefits, you can use AI to generate a comprehensive, persuasive set of reasons to believe. You can explore different categories of reasons to believe such as safe and comfortable to bolder and riskier or more innovative reasons that push boundaries.

Using LLMs to augment creative thinking or strategic thinking skills is important, too. So that’s like saying, “Alright, what product innovation might make sense here? How should I position the product? What plan might I devise to research this further?”

Leveraging these skills, we refined our approach to articulate the product concept to consumers more effectively. This was distilled into a succinct one-page template that encapsulates the product’s essence, using language that has been validated through early consumer testing.

By following the client’s testing protocols, we established a clear benchmark for what constitutes a successful concept – one that is likely to achieve high sales. Concepts developed with the support of AI not only met but exceeded these benchmarks with significant margins. Further, AI-supported concepts were more likely to exceed the benchmarks than those developed by humans alone. This strong performance, backed by solid evidence, garnered the full support and confidence of the leadership team.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Note on stakeholders and technology stack

DAVID: There's a whole range of different stakeholders involved through phases 1-4. There are the business people, who are more concerned about the project's objectives. Then there are the more insight and market research people, who are interested in obtaining the right audience insights. And finally, there are the product innovation or marketing people who will ultimately have to champion this and then market it.

From a technology stack perspective, the company created their own frontend onto OpenAI APIs – a RAG-plus model. There's a trade-off between RAG and a well-defined prompting process like ours where you bring in the brand positioning documents yourself and bring your insight. You could argue that RAG could probably do those things, but since it does them in a way that's opaque, you can never really be sure whether it's properly incorporated in the right context for you. So, you might as well just copy and paste what you need directly into the prompt.

A clever user, a deliberate user, should almost always pull that stuff in themselves rather than rely on the untransparent and unproven benefits of RAG. Certainly, I never saw any evidence that RAG was beneficial. We never depended on it in any of the processes we wrote.

Phase 5: Adoption and Integration

Our client is now engaged in the fifth phase, Adoption and Integration, which will be followed by the sixth phase, Evolution and Learning. We've conducted some training sessions, but it's now up to them to carry it forward. There will inevitably be an evolution as new prompting techniques emerge.

They could never have just jumped straight to phase 5. They had to go through phases 1, 2, 3, and 4 to get there. We guided them through these phases. It was a complex, messy project that evolved over time with shifting goals and targets, but we realized how crucial this was.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Results and Impact

DAVID: What started as assistance with AI ended with a total transformation of the innovation process. We trained hundreds of people and guided individual teams and countries worldwide on real projects, and this is now how they innovate. It was challenging work, primarily because the scope was continuously changing. As we and they learned by doing, we realized how transformative this process is.

In the end, it was a journey of about five months. I think it could have been shorter, but there was a massive number of stakeholders to include in our journey, and the scope of the project evolved significantly as we progressed. It began with a simple request for prompts and culminated in a complete redesign of their innovation process. If the brief had been to redesign the innovation process from the start, they probably wouldn't have approached a small consultancy like us and instead gone to a major firm like McKinsey.

I'm surprised that we didn't end up with a more definitive process. I think what we recognized both in the pre-AI and post-AI world is that every project is a little different and the different innovations require slightly different processes. There were instances when we began with piles of in-depth audience research, while at other times, we started with almost none. Sometimes, we excelled at certain parts of the project, and sometimes, we skimmed over them to focus on other areas.

What we ended up developing was a highly flexible and comprehensive process, encapsulated within a detailed 300-page manual that includes various examples. It's designed to allow users to navigate through it at their own pace. The impact on the company and its approach to product innovation has been significant. It's not just about the number of innovations, but also about their quality and success rate. It's a game-changer for the company, heralding a new era of product innovation powered by AI.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

Impacts included:

- A global initiative, with case studies in South America, Australia, Indonesia, and many other countries.
- The company has adopted the new manual for how product innovation should be done which they're in the process of rolling out.
- We proved that a human working with AI can generate more innovations that pass the benchmark score in consumer testing than a human working without AI, and score more highly when doing so. So, the company is getting more products that qualify as hits in consumer testing, which means a higher rate of product innovation success.
- We produced evidence that the partnership of AI and people is stronger than AI alone, which underscores the value of further investment in people who can empathize, delve deeper, and refine innovations that resonate on a human level and elevate results.

Key Learnings

- **AI-fueled creativity:** The combination of human ingenuity and AI is more potent than AI alone. It augments human potential, enhancing the analysis and exploration of large datasets. The human touch is crucial for asking the right questions at the outset and having original thoughts and connections.
- **Synthesize information and uncover gaps:** LLMs are used to synthesize existing information, process vast amounts of data, and uncover gaps in our thinking. This helps identify patterns and trends in data, including market research, customer feedback, and internal performance metrics, that might be missed.

GENERATIVE AI CREATES A PARADIGM SHIFT IN INNOVATION PROCESSES

- **LLM skill development:** The focus should be on assisting people in developing foundational skills to operate LLMs. The potential of LLMs in pattern recognition and generating innovative solutions is maximized when users are trained in its application. There are four key skill areas to master: condensing information, desk research, content generation, and creative thinking.
- **AI adoption across six phases:** The transformative journey of AI adoption requires active participation to fully understand and appreciate its impact. The six phases are: 1) Skepticism and hesitation, 2) Exposure and realization, 3) Experimentation and trial, 4) Realizing practical applications with benefit realization, 5) Adoption and integration, 6) Evolution and learning.
- **Change management initiatives:** The following nine initiatives can accelerate the process:
 - AI exploration sessions: Hands-on sessions for team members to explore tools.
 - Role-specific AI integration: Tailor tool use to specific roles.
 - New employee AI onboarding: Include tool training in the onboarding process.
 - Internal AI tool development: Encourage the development of company-specific tools.
 - AI success case studies: Share stories of successful implementations.
 - Regular AI tech talks: Discussions on the latest developments.
 - Cross-functional AI collaboration: Encourage cross-departmental use of tools.
 - AI mentorship circles: Pair AI-experienced employees with newcomers.
 - Tailored prompts for your organization: Customize prompts to fit the organizational context.



30 CASE STUDY

Unlocking Profit Potential – Leveraging Enterprise Data for Customer Profitability

IN CONVERSATION WITH BEN HALL: Director of Group Finance at maritime solutions provider V.Group. Prior to this, Ben was the Finance Director at Bupa, a leading global healthcare specialist. His extensive experience also includes senior roles at various financial services companies, and he began his career at Ernst & Young.

LinkedIn: [ben-thomas-hall](#)

Case Study in Summary

Case Study Classification

- Segment: B2B
- Industry: Ship management and marine services
- Geography: Global with 60 offices worldwide, including head office in London, UK
- Value creation levers: The utility and investment return of a service portfolio
- Time to impact: 12-18 months
- Regulatory considerations: Low

The Business Challenge

V.Group wanted a coherent understanding of customer profitability, which was compromised due to fragmented data, disconnected systems, and incomplete information. Furthermore, there was an over-reliance on compiling data in Excel to try to understand the business. This approach was not providing reliable and actionable insights, so they needed a more integrated solution.

The Response

The company's data management journey began with the development of a clear vision and roadmap and the engagement of business stakeholders. The company was leveraging data for ship operations but needed to focus on the internal use of information for better company management.

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

The work was done in three phases:

Phase 1: Developing a clear vision and roadmap, and engaging business stakeholders. This phase shifted focus from ongoing data manipulation to laying out a vision, ensuring buy-in from key stakeholders. The company culture, previously centered on ship operations, began to appreciate the internal use of information for better company management. The project was rethought to focus on the rule rather than the exceptions, which led to an effective revenue and cost allocation approach.

Phase 2: Implementing data management and incremental improvement. This phase prioritized data capture and management, with an emphasis on automation and alignment of data across systems to reduce manual effort and streamline processes. This involved identifying bottlenecks, seeking alternative approaches, and making incremental improvements.

Phase 3: Maturing and scaling. This phase evolved the data collation processes with the use of a business intelligence tool, predictive work, and dashboard provision. It set up the company to apply AI prediction to internal data in the future. The business will leverage the “one data platform” approach for self-service analysis and reporting.

Case Study in Detail

Ben Hall profile

Ben Hall is Director of Group Finance for V.Group. His role is focused on measuring, reporting, and influencing the financial performance of the organization, as well as providing a projection of performance.

V.Group is a large-scale operation with a wide portfolio of complementary services. The group employs just under 3,000 onshore staff and has an annual revenue of around \$700 million. It provides technical management, an array of services, and digital solutions for the running of large maritime vessels on behalf of owners. It encompasses manning services worldwide, including a pool of 45,000 seafarers, with 25,000 on board ships around the world at any given time. As the group doesn't own the ships, it's a capital-light and very people-reliant organization.

There are many sub-businesses within the core offerings, making it a one-stop shop for a vessel owner. It provides training for seafarers and offers a range of operational services, such as repairs, maintenance, procurement, insurance, travel, and catering.

The Business Challenge

BEN: My roles within this and previous organizations have been very focused on performance management where you're looking underneath the numbers and the financial outcome to what the drivers of the business are. To really understand the drivers, you have to take an interest in understanding the product and operations, whether that's working in an investment bank, an insurance company, or a maritime services provider.

You then realize that profit is just the end outcome. You have to take a step back and ask, "What are you going to measure and track and predict in the organization?" Ensure you're predicting and measuring the activity – the root cause of the outcomes – not the outcome itself.

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

You find yourself tapping into data from all parts, people, and systems in the organization. You very quickly realize that in large organizations there can be an industry of people who are consistently engaged in the task of interfacing with disparate systems, repopulating incomplete information, and then ultimately smashing it together in Excel, for example, to try and understand the business.

When I was brought into V.Group, I was tasked to look at performance measurement using financial and other pieces of data, with a focus on customer profitability.

Context-dependent issues

Dominant external focus

BEN: The company's strong customer focus with respect to data usage deflected focus away from internal performance metrics and management. This left a gap in understanding profitability across different customer segments. When I joined V.Group, owners and management would ask for information to track performance that they assumed was readily available. However, performance measurement relied upon financial systems structured to support external financial reporting only. Management information had taken a back seat.

Lack of proficiency in measuring impact

BEN: New company owners and management recognized that capture of the financial outcomes was focused heavily on individual legal entities. For example, the company did not have a ready response to key questions such as "How much does the company earn from a tanker versus container versus leisure vessel?" Service lines couldn't measure their impact beyond an approximation, let alone look across the cross-sell profitability of a customer.

A program had already started on profitability, but was losing traction

BEN: Consultants had already been brought in to build a model for customer profitability, but the lack of clear vision, direction, and management ownership meant their efforts were not progressing as expected. Without internal buy-in, the task had become absorbed in cleansing and manipulating data. There was a lot

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

of data and multiple systems that were disconnected and not managed in a way that allowed easy collation into one centralized view. Frankly, the consultants were very frustrated by the company's lack of support.

Approach

BEN: I have heard the term “you can’t boil the ocean” thrown at me when it comes to talking about data in the company. I realized that I was trying to lay out a vision and that it felt challenging in an environment where data quality and completeness were not a focus. So, of course, that vision needs to translate into a plan containing manageable phases, actions, and deliverables.

A three-phase approach was taken. The phases are not strictly linear and may overlap or be revisited as the project progresses as it’s still ongoing, with a focus on automating data capture and management as much as possible.

Phase 1: Developing a clear vision and roadmap, and engaging business stakeholders

BEN: You need to have conviction. In any role in any company, there’s a lot going on and there are lots of competing priorities. When trying to develop something that you know is the right thing for the future of the organization, the company may be fighting fires daily, but you need to find a way to keep the momentum going.

First of all, there had to be a clearer picture of what we were trying to achieve. The efforts to that point had been channeled into data manipulation versus taking a step back and laying out the vision. Without a clear vision, you don’t get the appropriate buy-in from key stakeholders, including operating management. Taking the time to think through the direction and to articulate it properly to the right people was crucial.

Even if you are very IT-reliant, you need business sponsors. Perhaps the company doesn’t have the skill set outside of IT, or in some cases, the right team design to work with this sort of program. My advice is to find ambassadors who can

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

see the vision and are fully bought in to drive it, keep efforts on track, and get management on board.

It sounds obvious, but when setting out to build a usable set of data to provide answers, it is helpful to understand the most relevant questions. We spent time thinking through business use cases, and I asked what was important to senior stakeholders to steer the shape of the data we were developing and how we could apply it.

I encouraged our consultants to steer efforts toward surfacing the data and exposing the project: “Let’s look at this in the 80/20 space. Rather than focusing on the exceptions, let’s focus on the rule.” We realized that we got to a point where we had good revenue allocation, even if it was not yet automated. But the allocation of cost was proving heavily manual and differentiated across services, while also blurred by accounting practices and system classifications.

V.Group does not have a uniform product, so cannot apply a simple or unique cost allocation. We have many differentiated customers and services, and they require a different level of attention, level of service, and effort, and can have entirely different business models. Keeping materiality in mind, we had to look at each of those different service lines and seek out the simplest and most appropriate cost allocation approach with the available drivers.

When you’re dealing with a decentralized and widespread environment, you may not be seeking 100% accuracy all the time. We aimed to be directionally correct and be about guidance rather than a perfect solution. It’s been quite a long road of working with the business leaders, their finance teams, business partners, and others to the point where they can see the opportunity in the data and the reason for it.

This phase had involved establishing a clear vision, setting a timeline, and working with consultants. We named the project “The Pod” (Profitability Optimization Data), and over time the acronym for the data has become a noun, referenced by users across the business. The name seemed to soften the finance-centric connotation of customer “profitability.”

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

Phase 2: Implementing data management and incremental improvement

BEN: This phase moved the process into the bigger picture of data capture and data management. Extensive manual effort including downloading and data manipulation isn't sustainable when you target regular and timely information. An alternative approach was needed. Over time, we achieved automated data retrieval and captured the mapping and alignment of this data across systems as much as possible.

We progress by continually checking in with practical questions: "What part of the process is still taking time? Why are we held up again this month? Now let's look at this piece. How do we make that smoother?" And you improve incrementally. This takes patience.

Again, a very high priority was building ambassadors in the business. Getting the information out there very early on earned an appreciation of what was being produced, which then allowed us the resources, or at least the patience from the organization, to start improving the process and making it more fluid and efficient.

Attributing revenue to a customer by following the invoicing is straightforward. However, the effort was in extracting that particular data from different systems and automating that data capture without having to run around and go into every system, every month, and manually download it. We now have a mechanism for extracting data from key systems or key modules within the systems on the revenue side.

However, on the cost side, the process can be more intensive, particularly when you are trying to allocate people and costs. This phase was about defining the cost categories and how to size that cost, defining the allocation key and how you capture that allocation key. At this point, I have found that the secret to success is simplicity not complexity, but that's not as easy as it sounds. A business leader will rarely accept being measured on an opaque, black box calculation, so breakthrough can occur through simple but appropriate allocation keys, such as headcount.

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

Having developed a single data platform, the company significantly evolved the infrastructure for data management across systems, which will be leveraged for all ongoing data applications, including customer profitability. Currently, data is extracted into Excel or Power BI for self-service analytics, visualization, and reporting; however, we are excited about future improvements.

Phase 3: Maturing and scaling

BEN: When the data was mature enough, we created simple dashboards and gave them to the business leaders so they could see the information live and drill down into details themselves.

It started with hands-on access but moved on to scorecards, performance management, and internal measurement. We now monitor the businesses not based on their business line but based on their customer remit. This places accountability with the customer relationship owner and breaks down silos between functions that cross-sell to those customers, something that many companies struggle with.

The company does apply AI prediction to support the performance of customer assets, but not for internal data, although it's definitely where we're heading now that we've taken a wider view of multi-system data management. Currently, our customer profitability has surfaced with the provision of dashboards and some element of drill-through and source data that can be accessed by analysts directly. We have migrated toward Power BI as a universal tool. This gives us an aligned way of self-service analyzing, reporting, and creating reports for sharing and monitoring across the organization.

Results and Impact

BEN: The executive team and company owners found the customer profitability model gave them an invaluable new perspective. It allowed them to view the company's financial performance in a way that points to the operational source.

The company changed the way it looked at, presented, and modeled itself. Instead of focusing on business lines, there was a shift to customers and vessels. This

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

gave us better alignment of performance to market opportunity. We now model our company and plan based on the customer type, rather than the business line first, which is what we did previously.

To bring this to life, you might consider the market opportunity for the tanker market, which has different dynamics to, for example, container ships. They operate with different supply and demand cycles, so by being able to view the business by customer and by vessel, we instantly started to ask, “How does that map against the external environment? How does that map against our view of the future environment?” We now critique our financial forecasts with this market lens.

In business operations, you want the application of this information to be a little bit different, more focused on the customer, the product offering, the pricing, and so on. In this interaction, you may face denial as data may expose weaknesses in the commerciality of a customer relationship. But, critically, it exposes opportunities.

I advocate getting information out there, in the hands of the users, as soon as possible. By building early advocacy, people start caring and it buys support to invest time and effort into improving the product. It also exposes early weaknesses, so you can course correct. Our allocation process has automation coded in; however, there are residual elements of Excel and manual reliance. We are pushing for a better, more systematic way of doing this. I would recommend seeking advice on data capture, management, and tooling from the start, rather than trying to retrofit.

We have been on this journey for years, but that’s not to say we didn’t have anything to apply throughout that journey – we absolutely did. We’re two or three years into having a robust, timely, monthly dataset. And before that, we probably had a year or two where we were quarterly. Even prior to that, we were seeing big-picture information not previously available.

Company-wide manual processes consume more resources than it seems on the surface. By streamlining processes and getting to the point where they’re predominantly automated, we have data for application in performance measurement

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

intra-month. Having timely, actionable, and fact-based information – that’s where you start getting value.

Impact included:

- The customer profitability model provided the executive team and business management with a new perspective and detailed understanding of the company’s performance.
- The business is now analyzed through both customer and vessel perspectives, allowing clearer mapping of performance and future projections.
- This data is widely applied for internal decision-making, holistic performance management, and to underpin a strategic focus on the quality of earnings.
- Data collation processes are continually evolving, with plans to migrate to more advanced systems and tools.
- The company applies AI prediction to customer vessel data, with plans to extend this to internal data.
- Automation and streamlined processes have resulted in timely performance measurement, adding value to the business.

Key Learnings

- **Maintain focus with a roadmap:** It’s crucial to have a clear vision of what you’re trying to achieve with a well-defined AI roadmap and be able to articulate it to the right people. Quantify the problem to solve and the gap between the current state and the proposed solution. Ensure the problem is manageable in scope for resolution. Keep the momentum going on your project, even when there are competing priorities.
- **Ask the right questions:** To extract the most meaningful output from the dataset, you need to ask the right questions. Without the right questions, the dataset will just be another breakdown of results. Formulating thoughtful questions upfront guides the analysis and transforms the data into actionable intelligence.

UNLOCKING PROFIT POTENTIAL – LEVERAGING ENTERPRISE DATA FOR CUSTOMER PROFITABILITY

The goal isn't only financial quick wins, but to understand the drivers of the business and use this to improve performance and decision-making.

- **Business sponsorship:** You need key stakeholders outside of IT who can see the vision and are fully bought in to drive efforts, keep them on track, and get management on board. Sharing information early in the process fosters appreciation for the project, enabling resource allocation and patience for process refinement.
- **Consider the 80/20 rule:** To get something out there, it's okay to release a partial end product. When other people start caring about it, you buy the headroom to fix things that couldn't be fixed during the initial build.
- **Build for scalability:** To avoid having to retrofit your architecture and tooling later, seek advice on the best ways to build your model and the right tools to use at the start. This involves lining up the systems for data capture and manipulation, and the tool to surface it more productively.
- **Be patient:** Understand that building a robust, timely, monthly dataset is an iterative process. You need to have patience as you build up a credible history.



31 CASE STUDY

Minimizing Customer Churn with AI

IN CONVERSATION WITH MANU KUMAR: Chief Data and AI Officer, with over a decade of experience in the advanced analytics and AI space. His expertise spans senior leadership roles in technology, luxury, eCommerce, and insurance sectors across three continents.

LinkedIn: [manuakumar](#)

Case Study in Summary

Case Study Classification

- Segment: B2C
- Industry: Insurance
- Geography: Europe
- Value creation levers: Growth and retention
- Time to impact: 12-18 months
- Regulatory considerations: High

The Business Challenge

A large, multinational health insurance company with revenue in the low double-digit billions of dollars and operating in a highly regulated industry was suffering from high churn rates across its major countries of operation, implying losses in the current year and future revenues. This intensified pressure on customer acquisition costs as well as inefficiencies in targeting marketing spend. The company was held back by a lack of data and AI solution culture and expertise and high switching costs from a legacy solution that was deeply integrated into actuarial processes like pricing and risk modeling. Any changes required regulatory and stakeholder approval.

The challenge was:

- Develop an effective solution to reduce churn.
- Prove the value of investments in AI capabilities and deliver a strong, quick ROI.

MINIMIZING CUSTOMER CHURN WITH AI

The Response

Manu was a key member of a new internal team called Customer Lab, created to overcome cultural resistance and build transformative capability to help the insurer modernize its processes.

The team focused on developing an AI-based churn prediction model using an expanded set of 40+ inputs from the insurer's existing data to identify patterns to predict future churn. These included potential drivers across categories such as product or policy factors (for example, premium increases, discounts, free cover duration), customer profile (age, tenure, and so on), claim patterns and complaints (number of claims, payment delays, and more), as well as signals related to adding dependents to an insurance policy (potentially making certain pricing unattractive for a policyholder). They trialed a variety of AI approaches and finally settled on a Long Short-Term Memory (LSTM) model – a deep neural network that can use time series information of events to make predictions.

The work was done in three phases:

Phase 1: Discovery and building the team. This phase focused on understanding the core business challenge and current approach and assembling a new team versed in technologies to be implemented in the solution.

Phase 2: Developing a minimum viable product (MVP). This phase involved a 12-week sprint to construct a data-driven MVP to predict customer churn, addressing shortcomings of the existing model.

Phase 3: Driving insights, operationalization, and adoption. This phase identified key churn drivers and decision matrices and devised customer retention strategies. It enhanced predictive performance accuracy, addressed risk concerns, and promoted widespread adoption.

This overall approach significantly improved the performance of the model and the ability of the company to proactively identify customers at risk of leaving the business. The model could be applied to 95% of those who had been customers for over a year and could accurately identify 80% of customers who were likely to churn in the 12 months going forward.

Case Study in Detail

Manu Kumar profile

Manu is currently the Chief Data Officer at Richemont Group's YNAP group (Net-a-Porter, YOOX, Mr Porter, and theoutnet.com), focusing on building AI solutions for brands and eCommerce. Previously, he held senior leadership positions in data and AI at an insurance business. Over the past ten years, Manu has focused on data and AI. Prior to this, he worked in product technology roles at companies including Microsoft, Vodafone, and various startups across three continents.

Manu says his diverse background gives him significant advantages. It leads him to take a “startup-like” approach, where, despite the size of the company, he enjoys getting hands-on, developing the strategy, and then coding alongside teams. In addition, he applies his experience from a variety of industries to relevant situations; for example, his Master's in Operations Research from Berkeley provided foundational knowledge in techniques like neural networks, which played a valuable role in predicting customer churn in this case study.

The Business Challenge

The insurer operating in a highly regulated industry was experiencing an expensive customer churn problem with low double-digit churn rates in its scale markets. This impeded revenue trajectory challenged realizing customer value against acquisition cost and exerted further pressure on the marketing and sales

MINIMIZING CUSTOMER CHURN WITH AI

organizations. At the time, there was no scalable mechanism in place to proactively identify high-risk customers likely to churn and take steps to prevent this.

Context-dependent issues

- **Lack of solution innovation culture and expertise:** The business, aware of data and AI advances, was seeking to modernize. Until that point, it had been good at sourcing externally built solutions and expertise via advisors but had not yet developed a strong internal innovation culture nor the ability to build tailored solutions for key strategic problems. At this relatively greenfield stage of its analytics and AI capabilities, the company set up a new customer lab team to accelerate its capabilities and hired functional specialists like Manu to prove value.
- **High switching costs:** The business's legacy solution was deeply integrated into core actuarial processes such as pricing, risk modeling, and regulatory reporting. There were significant switching costs as it needed approval not only by regulators but also from a wide range of stakeholders, including strategy, legal, compliance, actuarial modeling, and pricing.

Approach

MANU: When starting a project, I invariably draw on several frameworks, and in this instance referenced *Competing on Analytics* and *Analytics at Work* by Thomas Davenport and Jeanne Harris. While it might not be entirely relevant to new technologies, Davenport and Harris's DELTA model offers a valuable framework for understanding and implementing successful analytics programs. The DELTA Plus Model consists of seven key elements:

- Data: Quality, integration, and governance of data
- Enterprise orientation: Building an analytical culture and adopting enterprise-wide analytics

MINIMIZING CUSTOMER CHURN WITH AI

- Leadership: Strong executive support for analytics
- Targets: Applying analytics to the right business areas
- Analysts: Hiring and managing analysts effectively
- Tech stack: Having a tech stack that is as cloud-forward as possible
- Analytical techniques: Using appropriate methodologies for data analysis

I have used this framework across different industries from telco to healthcare and eCommerce. These factors are crucial in planning and implementing successful data science projects. With this in place, I focus on four key areas, based on my personal experience:

- Data capability assessment: Conduct a thorough and honest evaluation of your internal and external data capabilities.
- Modernizing models and tools: This is often the low-hanging fruit. By updating your models and tools, you can tap into opportunities that might currently be overlooked.
- Change management and transformation: Purchasing software and using algorithms are important but it's crucial to ensure your team understands the benefits of these tools, how to use them, and the improvements they can bring to their work. This often involves a significant change management effort.
- Talent: It might be hard to find but strong data science and data engineering talent is critical, as is having them in an integrated team. I've faced issues when data engineering was not part of the data science team. To deploy multiple models at multiple different points, a very strong data engineering team is absolutely key.

Phase 1: Discovery and building the team

MANU: I began by asking, “What’s the biggest business problem you have?” The resounding response was customer attrition. Acquiring new customers to replace your old ones is very challenging, so this gave me something meaningful to focus on.

The first thing I did was go on a discovery mission to figure out how the business was currently addressing this problem. Rather than reinventing the wheel, this meant I could identify a starting point from which we could build. It turned out that the actuarial team had a basic statistical solution in place as part of a wider actuarial software suite that created their prediction tables and pricing structures as well as tracking money in from premiums and out for claims. This was over a decade old and suffered from the impact of model and data drift. In addition, it depended on a dated, “small data” statistical approach that could not leverage the rich signals and datasets available to the company.

The key problem with this “solution” was that it had low overall accuracy due to two things. Firstly, the model was designed around a narrow set of inputs, using only a few predictor variables (age, gender, postcode, product type, claims, and premiums). It lacked the sophistication to leverage the insurer’s richer customer and claims dataset. Secondly, the simplistic modeling approach of the legacy solution meant it fell short of the business needs. The statistical, multivariate regression model it used didn’t consider the valuable time element of the customer’s journey. Some regression-based approaches can be useful, but in this instance, it was designed more as a simple “risk modeling” tool and not a purpose-built churn prediction tool that could provide predictive insights into specific time horizons. So, the business couldn’t use it to proactively mitigate churn risk.

MINIMIZING CUSTOMER CHURN WITH AI

In addition to low accuracy, the legacy solution was deeply integrated into existing actuarial processes such as pricing, risk modeling, and regulatory reporting. So, you can imagine the amount of resistance to anything new upsetting the apple cart because here I was walking in with a data science background, they thought, “How can he figure out our risks and exposure better than us?” Therefore, there had to be a high level of trust in a new solution. And there were significant switching costs to consider too – it would not only need approval from regulators but also from a wide range of stakeholders, including strategy, legal, compliance, actuarial modeling, and pricing.

We lacked data scientists, so I hired externally to bring in the necessary expertise with professionals who understood the technologies we needed to implement. They evaluated our existing algorithms, identified areas for improvement, and upgraded them to enhance performance.

Phase 2: Developing a minimum viable product (MVP)

MANU: We focused on developing a data-driven mechanism to predict churn and address key shortcomings of the existing model. I took in all the underlying datasets and created a new AI-driven system. This system used Long Short-Term Memory (LSTM) networks, which are precursors to ChatGPT. LSTM networks are widely used in deep learning and excel at capturing long-term dependencies, making them ideal for sequence prediction tasks. They can understand and predict patterns in sequential data like time series, text, and speech.

We outlined a 12-week approach to build and deliver an MVP (minimum viable product) that would demonstrate worth, gain traction, and start to earn trust.

MINIMIZING CUSTOMER CHURN WITH AI

Exploring datasets and inputs	Preparation: hypotheses and insights	Developing lapse prediction models
2 weeks	4 weeks	6 weeks
• Assess existing approach and workflow • Understand the scope and nature of input data • Engage with to establish success criterion	• Collect, clean, and aggregate relevant inputs • Present insights to business, consolidate hypotheses to test with data	• Define training and validation data • Detailed feature engineering to provide right inputs into model • Explore modelling techniques to predict lapses

Figure 31.1 – Twelve-week MVP process

Exploring datasets and inputs: As a first step, we undertook a scan of potential datasets that would help us understand the customer’s experience and likelihood of churn. Instead of the initial list of limited variables, we expanded the dataset to include more than 40 variables, a volume the previous in-house system couldn’t manage, with more signals than before. The original software, which had not been updated in a decade, exhibited significant data and model drift. This drift is a common issue in machine learning models where the statistical properties of the target variable, which the model is trying to predict, change over time in unforeseen ways. This can lead to a decrease in model accuracy.

The following table shows some of the categories we worked across. This richer dataset helped to significantly improve the model and provided extensive insights into the business. Given the sensitive nature of the data, all datasets underwent significant review and detailed processes, with the oversight of risk and regulatory teams to remove Personally Identifiable Information (PII) and maintain high levels of security.

MINIMIZING CUSTOMER CHURN WITH AI

POLICY OR PRODUCT FACTORS	CUSTOMER PROFILE FACTORS	DEPENDENT RELATED FACTORS	CLAIM AND COMPLAINT FACTORS	TIME SERIES
Number of products sold to customer Current premiums and excess Historical premium changes Levels of discounting Free premium periods offered Policy renewal timeframe Next payment timeframe	Age Gender Premium relative to other customers (deciles) Location (address) Household status	Number of dependents (partners, children, and so on) Claims made by dependents – value Claims made by dependents – details Claims made by dependents – count	Invoices and items generated Amount claimed Details of claims Number of claims Timeframe of claims processing Number of days in payments delay Number of customer experience issues (for example, complaints)	Time-stamp of claims and customer support Sequence

Table 31.1 – Categories integrated to enrich the dataset

MINIMIZING CUSTOMER CHURN WITH AI

Preparation: validating hypotheses and creating insights: Building a clean, consolidated dataset gave us access to key insights for the business to interrogate the data and drive proactive retention measures. This initial approach got the dataset into the hands of the business early and informed dashboards that a variety of teams could access and use to understand customers and behaviors. Examples of the analytics include:

- Volume of total customers and churners by business line, such as personal versus company insurance schemes.
- Mapping churners and customers to overall premiums.
- Average customer tenure.
- Where in the journey stage customers churn. We found that 60%+ of applicants lapsed at the time of policy renewal, and 50%+ contribute 80% of premiums.
- Compare the distribution of customers and average premiums against churn rates. This analysis helps identify patterns or trends, such as whether higher premiums lead to higher rates or the opposite, or if certain groups of customers are more likely to churn than others. It serves to make strategic decisions about pricing and customer retention.

Developing churn prediction models: We developed a sample of data to train our models. We used a simple funnel imager to help the business understand the chosen population.

MINIMIZING CUSTOMER CHURN WITH AI

Evaluating modeling techniques for churn prediction

With deep learning, I could go much further. We evaluated a number of AI techniques to model churn prediction. These included the Cox proportional hazards survival model, Generalized Linear Models (GLMs), and gradient boosting.

APPROACH	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Cox Proportional Hazards Survival Analysis	Developed by David Cox in 1972, this is a regression model. This model is particularly useful because it works for both quantitative predictor variables and categorical variables. Often used in churn prediction to predict time until an event (a lapse, in this instance) occurs. Commonly used in medical research.	Capable of handling data with varying time intervals allowing for flexible analysis.	Assumes that the risk (hazard) remains constant over time, which is not the case in forecasting lapses. Results can be complex for non-technical users to interpret.

MINIMIZING CUSTOMER CHURN WITH AI

APPROACH	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Generalized Linear Models (GLMs)	Developed by John Nelder and Robert Wedderburn as a way of unifying various statistical models, including linear regression, logistic regression, and Poisson regression. Serves as a flexible extension of ordinary linear regression models. Can handle both quantitative and categorical predictor variables, for example, count, binary outcomes, or rates. Used for applications such as predictive maintenance, insurance conversion modeling.	Ability to measure effect of input variables on churn. Offers flexibility to handle different types of data, including binary, count, and continuous outcomes. Results are easier to interpret.	Presumes a linear relationship between variables, which was not the case in predicting lapses. Prone to overfitting in specific segments. Less effective in capturing complex, non-linear relationships in data.

MINIMIZING CUSTOMER CHURN WITH AI

APPROACH	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Gradient Boosting	Developed by several contributors: Leo Breiman, Jerome H. Friedman, Llew Manson, Jonathan Baxter, and Marcus Frean. Machine learning technique used for regression and classification tasks. Creates a prediction model in the form of an ensemble of weak prediction models, typically simple decision trees. These models are added sequentially, each one attempting to correct the errors of its predecessor. The final prediction is made by summing up the predictions of all weak learners, each weighted based on its performance. Effective in churn prediction problems and risk modeling.	Ability to learn about relationship between independent variables on the dependent ones. Churn datasets are imbalanced by default (most customers stay); helped deal with bias in the dataset towards retained customers.	Was prone to overfitting. Challenging to use for learning lifetime behaviors and balancing weight of events over time. Can be computationally intensive.

MINIMIZING CUSTOMER CHURN WITH AI

APPROACH	DESCRIPTION	ADVANTAGES	DISADVANTAGES
Two-stage LSTM (Long Short-Term Memory)	LSTM networks are a type of recurrent neural network that excels in learning from sequences of data. Adept at applications involving time-series data.	Effective with sequential and time-based data and able to learn from both events and when they happened – making them valuable for complex churn problems. Ability to use rich transactional data at an account level.	Computationally resource-intensive. Complex to understand and explain in terms of specific variables to focus on and action to take.

Table 31.2 – AI techniques evaluated for churn

The LSTM model was our final choice. It offered us the advantage of identifying patterns not only from the “events” in the expanded set of variables but also from the timeline series. This included factors such as the timing of the event, its frequency, and its sequence.

Transitioning from one solution to another, particularly in the realm of time series analysis, can make a significant difference. Using the most advanced time series solutions available at any given time can greatly enhance the accuracy and efficiency of data analysis.

MINIMIZING CUSTOMER CHURN WITH AI

One of the key advantages of modern methods is the ability to perform feature engineering. This process involves creating new input features for machine learning from existing variables, which can improve model performance. However, if you opt for a commodity solution that merely organizes your data in a certain way, using a few variables, you may miss out on the full potential of your data analysis. This is especially true if you don't have the time to create a Business Analytics Solution (BAS) book – a comprehensive document detailing your data analysis processes and findings. Therefore, investing time in feature engineering and creating a BAS book can help you leverage the full potential of your data and yield more accurate and insightful results.

To find the optimal solution, I conducted thorough performance checks and model explanations. I explained and likened this process to a Google search, where the interaction between a company and its customers is through language. Just as you type a search word on Google and it predicts what you might be looking for, businesses can use similar predictive models to anticipate a customer's next action, such as a churn event or a product upgrade/downgrade. The most successful companies have implemented this architecture for highly visible products.

The beauty of using LSTMs lies in their ability to capture and analyze temporal patterns in data. We used LSTMs to calculate Customer Lifetime Value (CLTV) and churn rates, creating cohorts based on these values. This approach allows us to factor in acquisition costs in a detailed manner.

We can examine each customer individually and decide on the most effective intervention. For instance, for a customer who is highly likely to churn but has a high lifetime value, we might decide to have an insurance agent call them directly. This is an expensive intervention, so we would only use it for a small list of customers. This strategy ensures that resources are allocated efficiently, targeting customers who are most likely to bring value to the company in the long run.

MINIMIZING CUSTOMER CHURN WITH AI

The CLTV model we developed was more advanced than any previous system. By using LSTMs, we were able to leverage past transaction data for feature engineering and model development. Of course, all data was anonymized in compliance with the General Data Protection Regulation (GDPR). This approach allowed us to create distinct predictions for different cohorts, particularly regarding their likely claims as a group. The necessary data was readily available from past claim submissions.

The importance of time series analysis in this context cannot be overstated. It allows us to layer events into a timeline, providing a more comprehensive view of customer behavior.

The maturity of LSTMs played a crucial role in our success because they allowed us to fully unlock the potential of our datasets. The solution was tailor-made for our needs, which significantly enhanced its effectiveness. By leveraging the capabilities of LSTMs, we were able to extract valuable insights from our data and create robust predictive models.

Phase 3: Driving insights, operationalization, and adoption

MANU: Our analysis provided valuable insights into churn drivers and potential interventions for customer retention. For instance, accounts with promotional incentives exhibited a lower churn propensity.

In hindsight, getting a step change in predictive performance was the easy part because there was a good response to the results. However, the real challenge lay in overcoming the perception of risk, explaining the analytics and benefits to the broader business, and fostering adoption.

I often advise Chief Data Officers that if the reports, tools, or decision engines they create are not being used or implemented, then nothing they did matters. It's not just about creating the system; it's about making sure people are using it and processes are updated.

MINIMIZING CUSTOMER CHURN WITH AI

We implemented several strategies to drive adoption:

- **Showing the value:** We emphasized the value of improved prediction accuracy to the business, particularly in terms of customer retention through specific interventions. We quantified this value in terms of Return on Investment (ROI), taking into account the cost of interventions and potential gains.

To provide a more tangible understanding of the potential benefits and trade-offs of different customer retention strategies, we devised “trade-off matrices.” These matrices demonstrated the viability of specific initiatives, such as offering free cover, based on their adoption and retention percentages. This approach enabled us to effectively communicate the potential advantages and compromises associated with each strategy.

- **Translating:** We realized that there was a need to translate our work into clearer and more understandable terms for people unfamiliar with AI and bring the benefits to life through analogies and live events.
- **Explaining prediction:** People tend to think of prediction systems in binary terms, as simply being accurate or not. We spent considerable time explaining the confusion matrix in detail – false positives, false negatives, and so on – and what the value of reducing false negatives would be for the business. We used the analogy of COVID testing (topical at that time) and virus spread to explain the importance of precision to avoid false negatives. A false negative in COVID testing meant you didn’t pick someone who had the virus, and they went around spreading it, whereas a false positive might mean a person just got one more booster. With this explainer, people really got the whole concept of false negative and positive scores, and how the former was potentially more serious.
- **Audience-based access:** Not all information can be shared with everyone, so our approach to accessing data is to grant it based on privilege levels. For instance, regulators and actuaries have access to more comprehensive information due to the nature of their roles. This ensures that sensitive data is only accessible to those who need it for their specific functions, thereby maintaining data security and privacy.

MINIMIZING CUSTOMER CHURN WITH AI

- **Actuary copilot:** To address the pushback from various teams who might have felt threatened by this technology, I explained how the AI-driven solution could “augment” actuaries, as an AI copilot would make them more capable rather than replacing them. Because this one architecture could deliver better insights, it would help the teams price risk better and prevent the business from leaking value (such as leaking 1-2% fewer customers per year). Plus, it could solve more problems by layering in customer lifetime value, to design campaigns, and so on.
- **Industry examples:** By comparing the predictive nature of the architecture to how Google predicts search phrases, gave the teams reason to believe and build trust. Showing how one of the world’s leading companies has adopted such a framework demonstrated how their business could forecast whether a customer might upgrade, downgrade, or discontinue a product or service.
- **Getting close-up:** To drive change management, the team and I initiated live events with lots of webinars, roadshows, presentations, and different Executive Committee forums and so on to get people on board.

Results and Impact

MANU: The implementation of an AI-driven solution dramatically transformed the business operations. It achieved extensive coverage and applicability across the customer base, significantly outperforming the previous model and enhancing revenue generation. Moreover, it optimized internal processes, paving the way for future technological advancements.

Impacts included:

- The AI solution covered over 95% of the business and was applicable to most customers with a history exceeding one year.
- It could accurately identify 80% of customers likely to churn in the forthcoming 12 months.

MINIMIZING CUSTOMER CHURN WITH AI

- An eightfold performance improvement over the previous regression-based solution and yielded tens of millions in additional revenue over the period.
- The model streamlined multiple internal processes and enabled targeted, cost-effective retention campaigns based on finely segmented churn risk scoring.
- It unlocked additional use cases like predicting Customer Lifetime Value (CLV) and personalized pricing based on risk.
- The MVP's success gained approval for full-scale adoption to the company's largest market.
- The project established foundational data pipelines and AI deployment processes for future initiatives.

Key Learnings

- **Baseline technical capability:** Before introducing a new AI solution, assess the critical business problem to solve as well as the current systems (architecture, integration, performance, and scalability). This helps benchmark progress and identify opportunities, and risks. Understand foundational tech components to determine where to invest for risk mitigation. An important aspect of the assessment is to understand user workflows and the industry-standard technology that they rely on for critical processes. With due consideration, the business can integrate industry-standard solutions where appropriate.
- **Opportunities for quick wins:** Often, the initial outcome of a technical assessment identifies opportunities for quick wins. Even in an organization burdened with legacy systems and a suboptimal architecture, an AI solution can work effectively with minimal integration.
- **Identify in-house gaps in capabilities:** Assessing the capabilities of the internal team is also crucial at the start. Rapidly augment capabilities through targeted hires and strategic partnerships with specialist firms.

MINIMIZING CUSTOMER CHURN WITH AI

- **Data expansion:** Leveraging a broader set of data inputs can significantly improve the accuracy of predictive models and provide a more comprehensive understanding of customer behavior and potential churn indicators.
- **Algorithm selection:** The primary goal should be to solve the problem at hand, choosing and combining algorithms based on their suitability for the problem, not convenience. Understand the pros and cons of each method or algorithm. Update select algorithms and tools to offer immediate improvements in the development of AI solutions.
- **Minimum Viable Product (MVP):** Hands-on rapid development of an MVP provides quick progress within a typical 12-week delivery timeline. This includes bottom-up use case preparation, design and UI, build, UAT/QA, and deployment.
- **Trade-off matrices:** Effectively explaining precision versus accuracy and associated costs enables informed trade-off discussions. Developing “trade-off matrices” can provide a tangible understanding of the potential benefits and trade-offs of different customer retention strategies.
- **Buy-in from skilled employees:** It’s crucial to ensure the work is understandable to colleagues less familiar with AI. Integrating the AI-driven solution as an “augmentation” to skilled employees for specific process steps can overcome resistance. Progress can be slower than anticipated as it can take time to build trust and buy-in. An iterative and user/customer-led approach that integrates feedback is essential to the ongoing process.
- **Strategies for solution adoption:** Using various strategies, including analogies, other industry examples, and live events, can drive change management and ensure the successful adoption of the new solution into business workflows. Setting up a separate team within the business to accelerate the solution by building expertise and a change-oriented mindset can enable the core business operations to continue uninterrupted.

MINIMIZING CUSTOMER CHURN WITH AI

- **Data security and access controls:** When working with sensitive data, ensuring security and compliance is vital, especially when handling Personally Identifiable Information (PII). Implementing an audience-based access control system can ensure only authorized individuals have access to sensitive data. This ensures ongoing data security and privacy.



32 CASE STUDY

Enhancing Marketing Strategies Through the Power of LLMs

IN CONVERSATION WITH DAVID BOYLE: Director at advisory firm Audience Strategies, experienced insights leader, and co-author of the *PROMPT* book series on using language models like ChatGPT to guide marketing and brand growth professionals.

LinkedIn: [beglen](#)

Case Study in Summary

Case Study Classification

- Segment: B2C
- Industry: Media and entertainment
- Geography: Europe
- Value creation levers: Growth and retention
- Time to impact: Less than 6 months
- Regulatory considerations: Low

The Business Challenge

A music events company needed an audience strategy to promote a live event to a broad fanbase for a famous DJ with a genre-spanning career of over 20 years. The challenge included:

- Identifying distinct target segments for the artist with such broad appeal.
- Limited budget meant conventional audience analysis was not possible.
- The expectation for automation without an understanding of the complexities involved.
- Unclear directives and tight deadlines.

The Response

In a proprietary approach, David leveraged ChatGPT, alongside his and the promoter's extensive skills and experience, to formulate an audience strategy in approximately 20 hours. This process would typically take 5-6 weeks using traditional methods.

The work was done in three phases:

Phase 1: Understanding the person's key dimensions of fame. This phase developed a deep understanding of the artist's key dimensions of fame and how they resonate with different audience segments. ChatGPT was used to synthesize decades of rich online press coverage and commentary into data points by identifying implicit themes not explicitly stated in the texts. It's a collaborative and highly iterative process that leverages the strengths of AI while also requiring human decision-making and expertise.

Phase 2: Defining audience segments. This phase defined audience segments based on various factors, including their needs, lifestyle, interests, and behaviors. ChatGPT was used as a tool to understand audience motivations and behaviors, solve complex problems, and empathize with the audience. Then, through a series of prompts, profiles were developed to describe the typical lifestyle attitudes, behaviors, and interests of likely fans. This process helped narrow down the original 15 drivers to 7-8 core targetable audience clusters.

Phase 3: Developing tailored positioning and strategy. ChatGPT was used to develop needs, pain points, and ideas that align with the values of each group. This helped tailor marketing approaches to engage each fan group. Throughout the process, human expertise was applied to refine prompts and curate viable outputs. This combined expertise from both the promoter and the analyst ensured practical and broad-focused strategy recommendations.

Case Study in Detail

David Boyle profile

In an accomplished career spanning over 20 years in marketing and insights, David Boyle has held senior roles in prominent media and entertainment companies, including BBC, HarperCollins Publishers, and EMI Music. His leadership in analytics and insights teams has enabled his clients to better understand their audiences to optimize and deliver targeted content, experiences, and interactions via marketing. Currently, he works as an independent consultant and author, as well as a speaker on leveraging audience intelligence for business growth.

David's core passion is for democratizing access to customer insights from data. In his professional experience, only the biggest projects get good treatment, in terms of resources for analytics. He aims to use advanced analytics and AI to help organizations develop tools to provide richer insights for smaller projects and artists. Today, he applies his learnings from Fortune 500 companies to advise small startups with limited budgets on how they can unlock the power of data to drive better decision-making.

While not a technical data scientist, David's strong industry and strategy experience has enabled him to develop a specialist skill set using Large Language Models (LLMs) to uncover niche, exciting ideas. He and his team have worked extensively with LLMs (mainly GPT), and he's authored five books for a variety of industries, guiding them on how they can leverage this technology effectively.

The Business Challenge

As one of the UK's largest music events promoters, the client's business orchestrates an extensive array of concerts, festivals, and club nights by identifying potential talent among musicians and strategically pairing them with suitable venues. They then handle every aspect of event production, marketing, and ticket sales to deliver high-quality audience experiences and optimize tour profitability.

ENHANCING MARKETING STRATEGIES THROUGH THE POWER OF LLMS

The company was organizing a live event for a multi-faceted DJ with a long career, genre-spanning musical appeal, and humorous spoken word element.

The diverse characteristics of the DJ (funny, cool, provocative) led to a lack of clarity on positioning for the promotion team who struggled to arrive at a coherent audience strategy. Ambiguity about which audience segments to focus marketing and promotions on risked diluting the overall effectiveness of the marketing. The business needed to develop a clear audience strategy and direction to ensure tour profitability.

Context-dependent issues

Lack of clarity invariably leads to low or no budgets

DAVID: In the creative industry, there are a lot of permissions needed from different stakeholders to drive a coherent campaign. When there's no alignment on approach, then no one wants to spend money.

In my experience, clarity brings confidence, secures permissions, and unlocks budgets. Being very clear and having evidence makes everything move forward smoothly, rapidly, and more effectively.

I think a lot of companies encounter the same problems as this promoter: a lack of clarity means the wider organization does not buy into the idea (or value) of audience research to help marketing be more effective, and consequently, there is insufficient budget.

In such situations, teams frequently wing it – implementing intuition-informed initiatives. While not wholly bad, this can miss valuable insights that a more detailed audience analysis project might reveal.

LLMs misconceptions and reservations

DAVID: Many people misunderstand LLMS, thinking they're just next-word predictors or are incapable of producing something original because they rely on existing data. This notion goes against users' desire for unique, tailored ideas and nuanced insights. Others, meanwhile, believe they grasp the tech but one bad experience can lead them to dismiss it. Another fallacy is the notion that the whole process can be automated with the touch of a button, but there's a ton of skill required, which is not automatable.

We tackled these flawed assumptions by emphasizing that effective results require expertise and iteration. When the project sponsor ran her own prompts with sub-par results, we had to demonstrate to her, the leadership team, the staff, and the artist himself the need for deep domain knowledge and the right prompting approach.

Knowing how to use ChatGPT and then how to iterate, prune, and refine at every single step using its memory is an innately human process. You've got to edit and own the output. You're responsible, not ChatGPT.

Approach

DAVID: The Marketing Director wanted audiences to be at the heart of the work for this event and asked me to support the company to develop a strategy. I saw this as an opportunity to showcase the potential of LLMS like ChatGPT to rapidly develop insightful audience segmentations.

When meticulously prompted with the input of expert collaborators and refined over multiple iterations, ChatGPT can help you understand your audience, innovate, and set a strategy.

To fast-track actionable insights, a three-phase process was applied to leverage generative AI:




Figure 32.1 – Three-phase process to leverage generative AI in developing marketing strategies (Source: Prompt.mba)

Phase 1: Understanding the person’s key dimensions of fame

DAVID: First, we used ChatGPT to comprehensively analyze numerous articles on the artist from mainstream media, music press, and fan sites over his career. With several different phases in his career and many aspects to his character, this process yielded profiles for audiences of different ages and music genres. By providing very specific instructions, the LLM extracted and condensed the most significant themes explaining the nuances behind the artist’s evolving public image and multifaceted fame.

In this process, we copied the text of the articles into ChatGPT as part of the prompting, providing it with new context and requesting it to extract the essence of that context to identify the underlying themes. This is something ChatGPT is really, really good at – identifying implicit themes that are not explicitly stated in the text. It’s a huge advantage that researchers can leverage at scale. It allowed decades of rich press coverage and commentary to be efficiently synthesized into data points on why different fan segments are drawn to this particular persona and musical output.

The prompt here is crucial – you have to be specific. Writing “Summarize this article” is an absolutely terrible prompt. So, I would say, “I’m interested in marketing a tour by this person. I’m looking to understand the deep underlying needs of his audiences. Use this article to come up with a list of as many deep and underlying needs that drive people to want to go see this artist as possible.”

We learned that in some instances, the more context you paste in, the less the model pays attention to that context. Prompting needs a lot of nuance about what you are looking for, for example, underlying dimensions of their fame

ENHANCING MARKETING STRATEGIES THROUGH THE POWER OF LLMS

versus reasons why they’re famous versus different types of things they’ve done in their career versus reasons why people might be excited about going to see them live. Each is a different set of wordings for prompts and provides results highlighting different angles.

Every stage demands human decision-making, following certain results and abandoning others. It’s a collaborative process with hours of iteration, discussion, and debate. As the researcher, I have reasonable knowledge of this person, of the marketplace, and of audiences. But the decision-maker is the marketing director. She knows this person and the audience better than me and she knows what’s useful. So, there was a lot of back and forth with her to say, “Hey, I think this list of themes is pretty good, here’s 20, and then here’s the 40 that didn’t make the cut.” Then we talk them through together to refine them further.




Figure 32.2 –Improving prompts in a step-by-step manner to achieve more effective outcomes (Source: Prompt.mba)

Prompting: A ten-ingredient recipe

1.  **Act as:** The role you want ChatGPT to assume. Professional (for example, a doctor, lawyer, or scientist) or fictional (for example, Sherlock Holmes or Cinderella). Sets context and tone.
2.  **I am:** Characteristics that will help ChatGPT tailor its response to your needs.
3.  **Context:** Helps ChatGPT understand the nuances of your situation.
4.  **My ultimate goal:** Your long-term objective or the desired outcome you want. Shows the bigger picture and respond in a way that supports your overall aspirations.
5.  **My proximate goal:** The short-term or immediate objective you're focusing on.
6.  **Constraints:** Could include time, budget, or other factors that impact the response.
7.  **My audience is:** Helps ChatGPT create content that resonates with your audience.
8.  **Your task is:** Be specific and concise.
9.  **Response style:** The tone/style you want ChatGPT to adopt. For example, formal, informal.
10.  **Response format:** For example, a list, paragraph, bullet points, or a step-by-step guide.

Figure 32.3 – Ten-point prompting recipe guide (Source: Prompt.mba)

Phase 2: Defining audience segments

DAVID: Through the multiple iterations and curation of this analysis, we got to 15 independent drivers of somebody wanting to go see this show. But drivers are not enough; you have to bring the audience to life. So, for each driver, we asked what kind of person has this need? What are their hopes and dreams? What lifestyle do they live? What TV shows do they watch? What magazines do they read? What other music do they listen to? This is crucial because you are not marketing to the need; you're marketing to the broader person in the context of

their life so their related needs are really important to develop segments. We go from needs to audiences, which, in turn, give us audience segments, in a further two or three prompts.

I queried ChatGPT, “You deeply understand audience motivations and behaviors. You have high business acumen, complex problem-solving skills, audience empathy, and creative writing skills.” Then I asked it to perform as:

- An analyst at a top-tier audience strategy consulting firm
- A strategic and business-focused person
- An ethnographer
- A cooperative researcher
- A creative writer

We used several series of prompts to expand and dig deeper, as described above, on each driver or facet of fame and develop profiles outlining the typical lifestyle attitudes, behaviors, and interests of likely fans who would be drawn by that specific aspect of the artist’s offering. This helped us boil down the original 15 to a list of 7-8 core targetable audience clusters, which is a manageable number to debate. None of these were characterized by age or music genre taste, but by deep underlying needs, which is a totally different way of thinking about marketing the event.

Previously, we would go through this laborious process ourselves, reading extensively and conducting interviews with a wide variety of people in the industry. It could take weeks to compile a useful list of dimensions for a topic or theme. In this instance, we did it in hours.

In the end, every one of these segments would be valid candidates to go after and position the event and marketing around, but you can’t do them all. So, you probably have to pick one, two, or three and then go deep and market to them. At the end of the day, this takes an awful lot of judgment. Because they’re all valid ways to describe people, you can’t really be wrong about these segments. The kinds of tests you can therefore apply are, for example, asking, “Is this segment

useful? Does it fit the creative vision of the person? Is it actionable?” Ultimately, it’s about instinct and this is one of the reasons that you can’t automate this stuff, because I can’t articulate all the criteria I’m judging it on – it’s a task that lives outside of ChatGPT.

Phase 3: Developing a tailored positioning and strategy

DAVID: Finally, we leveraged ChatGPT to develop needs, pain points, and ideas aligned to the values of each group, helping us begin to tailor marketing approaches to the values and preferences of each fan group. Asking ChatGPT to work out what creative writing, positioning on adverts and marketing, as well as the imagery that would engage and excite each group.

Throughout this process, we continuously applied human expertise by iteratively refining the prompts and thoughtfully curating the most viable outputs. This infusion of real-world business judgment prevented impractical suggestions or those with a narrow focus from passing into the final strategy recommendations. It was about combining my expertise, because I know what a good audience segmentation and strategy looks like, and that of the promoter, who knows the artist and their goals. I couldn’t have done it without the promoter, and they couldn’t have done it without me – and neither of us could have done it without ChatGPT.

To take a more detailed view, this is what the prompt process looks like:

Because ChatGPT works best with relatively small specific tasks, I asked, for each of the six segments, “Please write three bullet points on how we might reduce their pain points as it relates to a night out like this.” Then I asked for three innovations we could do to make the event experience uniquely meet the needs of this audience.

Then, for each segment, I asked ChatGPT to think through three positioning strap lines and ways to position the event that might uniquely meet the needs of this segment. I use a lot of prompts. This is the kind of thing we would chain and do automatically because we’re just thinking through stuff. We’re not going to use all of it; it’s more like brainstorming.

ENHANCING MARKETING STRATEGIES THROUGH THE POWER OF LLMS

The final result is a creative brief with key messages, tone of voice, and creative direction that would engage each segment.

Results and Impact

DAVID: The approach resulted in increased revenue (ticket sales), cost savings, and time to insight to drive the strategy.

The final cost of our innovative method was less than 5% of the usual expense for an audience strategy project, delivered in 20 hours instead of 5-6 weeks. This demonstrates the potential of large language models like ChatGPT in significantly reducing both time and cost in strategic planning, resulting in a more impactful event with a higher return on investment (in ticket sales).

It's very rare for a promoter like this to have a metric they want to move before doing the work. Here, the key metric was sales. The challenge is that there's no counterfactual to this specifically: how much would you have sold if we didn't do the marketing in this way? There was no A/B testing of marketing strategies.

Impacts included:

- The AI-assisted approach provided rich audience insights in just 20 hours over a few weeks, comparable to the £30,000+ of traditional qualitative interviews, surveys, and data analysis.
- This automated approach was executed at less than 5% of the usual budget.
- The artist's fanbase was clearly defined into eight clusters, which helped align the promoter's team and built confidence in the audience(s) to target.
- This led to tailored positioning strategies incorporating messaging, venue design, and event format to resonate with each fan subgroup.
- There was a strong consensus that the three-phase process led to an impactful and higher ROI event than would have been delivered otherwise.

- The promoter moved from a generic, mass marketing approach, relying on gut feel and assumptions, to an AI-enabled, more personalized audience strategy for this event.

Key Learnings

- **Perspective broadening:** Beyond augmenting skills and saving time, LLMs are important tools to bring in broader perspectives for business idea generation. They can generate a higher volume of ideas, but with a greater variance of quality. Therefore, this necessitates the human expertise to discern the most promising concepts. Importantly, LLMs don't replace human creativity in the crucial task of combining these ideas and parameters into innovative, high-quality solutions.
- **LLM prompt optimization:** LLMs like ChatGPT are augmenting business processes by speeding up operations, providing valuable insights, and aiding strategic decision-making. However, ChatGPT requires specific detailed prompts to achieve optimal results. Furthermore, the process involves multiple steps and iterative refinement with expert curation at each stage to ensure high-quality outputs and prevent error propagation.
- **LLMs democratize data and analytics:** LLMs provide smaller businesses access to data and analytics tools that were previously exclusive to larger corporations. LLMs provide an affordable alternative to traditional market research and audience segmentation, leveling the playing field and facilitating growth and innovation across all business sizes.
- **Human-AI collaboration:** While automation can produce satisfactory results, excellence requires human decision-making. The effective use of AI demands human involvement at each stage of results, to optimize quality and prevent error propagation. This expertise and judgment cannot be automated.

ENHANCING MARKETING STRATEGIES THROUGH THE POWER OF LLMS

- **Industry-savvy decision-makers:** Collaboration with decision-makers who have a deep understanding of the industry, customer, or product is crucial. Addressing skepticism and knowledge gaps, you can ensure strategies align with their vision and goals.
- **Continuous learning:** Effective use of AI requires both industry knowledge and continuous learning. Even with guidance and training, there can still be gaps in understanding how to effectively use AI tools, which makes upskilling an ongoing imperative.



33 CASE STUDY

Multimodal LLMs Redefine Software Development and Customer Innovation

IN CONVERSATION WITH ZEEV FARBMAN: Co-Founder and CEO of Lightricks, an award-winning app developer that creates easy-to-use apps for photo and video editing. Lightricks's mission is to bridge the gap between imagination and creation, by building innovative photo and video tools that enable content creators and brands to produce engaging, high-performing content.

LinkedIn: [zeev-farbman-b5926b37](#)

Case Study in Summary

Case Study Classification

- Segment: B2C primary and B2B secondary
- Industry: Software
- Geography: Global, headquartered in the Middle East, with offices in the UK and US
- Value creation levers: Software development
- Time to impact: 6 months after initial consumer AI release
- Regulatory considerations: Low

The Business Challenge

With AI driving rapid changes in traditional business models, software development company Lightricks wanted to respond to new customer demands and future-proof their growth, by looking for more immersive, personalized experiences to drive consumer engagement. The maturation of LLMs and the development of Multimodal Large Language Models (MLLMs) have empowered AI to not only generate high-quality code but also automate complex tasks. This shift could commodify laboriously developed skills and processes and meet customers' changing expectations.

The Response

Zeev Farbman returned to his engineering roots to grasp the transformative impact of AI and LLM technology for his company. This technology's ability to process and transform photos and videos paves the way for automating tasks once solely in the human domain. Recognizing that it could streamline complex skills built over time, Farbman understood the need to respond and adapt.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Farbman shifted his business internally and externally, focusing on understanding how customers were integrating AI into their workflows and improving the speed and agility of his own business processes.

This case study demonstrates the profound impact generative AI will have on software development, and tools for creators and brands to enhance their creative process, opening up new opportunities for innovation.

The work was done in three phases:

Phase 1: Engage. This phase involved a deep understanding of AI technologies and their challenges to traditional operating models and internal processes, as well as their impact on customer segments.

Phase 2: Adapt. This phase involved making necessary adjustments to the operating model, organizing into relevant customer verticals, building new tools based on customer trends, and integrating AI into internal processes.

Phase 3: Innovate. This phase included exploring new ways to use AI to improve business processes and create new opportunities for growth through iteration and relevant application into customer products and services.

Case Study in Detail

Zeev Farbman profile

Entrepreneur, CEO, and Lightricks co-founder Zeev received a PhD in Computer Science, writing a paper on diffusion maps for edge-aware image editing.

He co-founded Lightricks in Jerusalem, Israel, in 2013 with Nir Pochter, Yaron Inger, Amit Goldstein, and Itai Tsiddon. The founders noticed that Adobe was not prioritizing developing consumer content creation tools for mobile. Seeing this as an opportunity, they took on the challenge to be “the Adobe of mobile” – a successful business strategy.

MULTIMODAL LLMs REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Today, the company's HQ can be found on the Hebrew University campus, in addition to offices in Haifa and Beersheba, and global sites in London and Chicago, which, in total, employ approximately 600 people.

Lightricks offers a range of apps to creators of all skill levels, including Facetune, Videoleap, and Photoleap, all designed to make it easier for users to express themselves and tap into the magic of creating. Its apps have been downloaded more than 680 million times.

The Business Challenge

ZEEV: Historically, Adobe focused on desktop and built 2D tools for professionals, while their competitor Autodesk focused on 3D tools for professionals on desktop. We saw an opportunity to prioritize developing consumer content creation tools for mobile.

I believe AI is going to have a radical impact on society and how we do things, and that traditional domain boundaries are going to change. LLMs have advanced to the point where they can generate sophisticated, high-quality code. This is a shift toward a new computing paradigm where complex human ideas and nuanced natural language can be converted into formats machines can understand. Such a transformation has the potential to revolutionize the software development industry – and we need to keep pace.

Customer and consumer expectations are shifting fast and there is more demand for immersive and personalized engagement, as seen around the release of a recent fantasy heist action movie. A global entertainment company came to us to explore a unique new way to market their new fantasy movie while emphasizing user engagement through technology, with a custom AI avatars feature.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Context-dependent issues

A new paradigm of computing

ZEEV: The complexity of technology, process, and data environments that organizations have invested in over time has resulted in a technical debt that makes it difficult for them to adapt and innovate. However, AI has the potential to revolutionize the way businesses operate and create new opportunities for growth and innovation.

Traditionally, computer code has been written in a programming language that machines can understand. It's an approach that needs a high level of technical expertise so it's limited because not everyone can do this. But natural language processing is a real game-changer because it can create high-level, high-quality code that's more accessible to everyone. It's this new paradigm of computing that has the potential to revolutionize the way we interact with machines.

An increasing number of projects where code is written by AI

ZEEV: The biggest disruption from our perspective as a software developer is that LLMs are already getting to a maturity where they can create high-level, high-quality code. We are starting to see more and more projects inside our company where most of the code is written by AI, which is leading to a new paradigm of computing.

This is evident in the development of a new breed of language models – Multimodal Large Language Models (MLLMs) – that can translate high-level human intentions and fuzzy human language into something machines can understand. They can combine text with other forms of information, such as videos, audio, and other sensory data, to understand and manipulate visual elements, solve problems, and unlock new applications.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

This new wave of technology has the potential to make many software development capabilities into commodities. For example, MLLMs can automate image and video captioning, speech recognition, and natural language processing, tasks that previously required human intervention to develop interfaces. If we don't adapt and build on this new technology, a significant portion of our work could become obsolete.

Approach

ZEEV: It's my firm belief that to stay ahead in business, we must continuously engage, adapt, and innovate, leveraging agile methodologies and technology to push boundaries and unlock new opportunities for growth. This is what we've done at Lightricks.

We implemented a three-phase approach to transform our internal processes and meet new customer demands. This involved a systematic and structured progression, using AI internally within Lightricks and externally to innovate with customers.

Phase 1: Engage

Using AI internally within Lightricks

Familiarize yourself with AI

ZEEV: The guidance I would give to CEOs and boards is that it's crucial to understand these technologies in depth. Often, the only reason things are done a certain way is due to inertia, but it's important to challenge the status quo. So, dive in and don't be afraid to ask those hard questions because it's the only way to ensure you're making the most of these incredible advancements in AI.

With this in mind, I decided to get back into programming to understand these technologies more intimately. This had a profound impact on our business, changing how we invest our resources, prioritize projects, and approach our work.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Assess the potential impact of AI on your business model

ZEEV: AI is so disruptive, especially in software development, that it's crucial to understand its potential impact on your business.

In the realm of online text, databases have become a commodity. Since Steve Jobs introduced the Macintosh in 1984, companies have invested heavily in building user interfaces. In fact, what many tech companies have been selling for the last 40 years is essentially layers of user interface on commodity technologies.

We're already seeing that AI can create personalized interfaces and, in the software world, we're moving toward having an AI agent that can build personalized software. So, we need to ask, if AI can build user interfaces and these interfaces become a commodity, what exactly are software companies selling? This level of disruption requires a close understanding of the incoming wave. It's essential for CEOs and boards to get close to this incoming wave.

Once you have a good understanding of AI and its potential impact on your business, it's time to explore the potential for both customer products/solutions as well as internal processes.

Using AI externally to innovate with customers

Understand the impact of AI on your primary customer segments

ZEEV: In the consumer segment, people are excited to use the latest tools. We conducted a survey earlier this year in the US sampling about 1,000 creators to understand their receptiveness to AI tools and businesses in general. We had the assumption that creators were using AI but were surprised to find that so many were already using AI in their workflows.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Key insights include:

- 84% of current creators are likely to use AI if it saves time or money, while an even higher percentage, 86%, are likely to use it if it positively impacts the creative process.
- While some creators have embraced AI as a tool for artistic expression – 64% of current creators consider it to be a form of art – others may still be skeptical or unsure of its potential, so it's still evolving.
- As generative AI becomes more accessible, we predict that the creator economy will undergo a widespread boom as creators can use it to scale their creative process and attract more brand deals and partnerships.

Phase 2: Adapt

Using AI internally within Lightricks

Introducing flexible engineering

ZEEV: Just before the generative AI wave hit, we diversified into four different business verticals to give our products and customers specialized attention. Three of these verticals are focused on photo and video editing apps that together serve millions of users. This split has led us to easily adapt and adjust our offering to include state-of-the-art AI capabilities.

The fourth vertical is more of a B2B focus and is based on a platform called Popular Pays, which we acquired in 2022. The platform connects creators to brands to engage them for content creation and quick, social-media-driven campaigns. In our B2B business, we're in direct contact with representatives from a number of brands responsible for creating content and marketing materials. For example, a brand might invest in a commercial marketing campaign only to realize that everyone has moved to TikTok. So, now they need to adapt their materials for TikTok, which requires vertical videos rather than the short, horizontal format in high-end commercials. AI can step in here and solve these problems quickly and efficiently. Even product photography, which can be costly and time-consuming

MULTIMODAL LLMs REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

to adapt to different geographies, can benefit from AI solutions. There's a big incentive for brands to start looking at AI solutions across the board.

It's important to note that brands have different concerns from consumers. They care about the legal aspects of using these models and where the liability lies, for example, if artists claim that their photos were used for training the models. However, things have recently changed dramatically as companies like Adobe have decided to take the liability of using these tools on themselves, reducing perceived risk.

Using AI for internal processes

ZEEV: We're embracing AI not just in our product offerings but across our internal operations as well, and the benefits are becoming clear. I believe it's essential for businesses like ours to adapt; those who don't are at risk of falling behind as technology reshapes roles across development, marketing, and design. If you're not part of the game, you're likely to be out of it pretty soon.

We're actively refining our approach to coding using LLMs and have developed a bunch of internal tools for different departments within the company. Adapting our processes to integrate these innovations presents its own hurdles, particularly when getting buy-in from our teams, even those who are technically minded, yet our commitment to transparency, especially when showing our new ways of working to clients, helps demystify how we've developed a different response to their marketing question.

Gradual adoption

ZEEV: Historically, computer scientists have made things more efficient, and now our jobs are online. Initially, there was excitement from the CTO and others, but this didn't immediately spread throughout the organization. Early attempts to deliver AI-based features to customers weren't all successful, leading to some internal backlash. This created friction with the teams responsible for implementation.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

We realized that such a significant change needs to be done gradually, with the entire team. So, we invested in ensuring everyone had access to the latest tools and found early adopters to work on challenging and interesting products. This required effort as the legal team had to develop guidelines on what's acceptable to use for coders, marketing people, and the IT team.

We started working with OpenAI early on and were one of the first companies they rolled their enterprise features to. We ensured every employee could use the latest models for language and image generation. Gradually, we started to see a positive impact.

A small cybersecurity startup I recently spoke with claimed that 95% of their first product was built using LLMs. This suggests significant gains in developer productivity. We used to believe there was no silver bullet for software engineering, but now we're realizing that we might have found one in AI.

Using AI externally to innovate with customers

Consumer trends and building new tools for recurring revenues

ZEEV: In the consumer sector, content creation is driven by trends. What people are seeing on social networks often directs the work. We've noticed hobbyist communities creating really cool things, and when people see them on Instagram, they expect to see something similar in the content creation tools they're using.

For example, there was a trend this year that was a huge hit for us. It allowed people to reimagine real-world locations as apocalyptic scenarios, inspired by apocalyptic movies and TV shows. We created a feature – AI Scenes – that made it easy to do this, and it led to a huge surge in downloads. At the peak of the trend, we were getting half a million downloads a day and were one of the top five apps in Apple's App Store.

But it's not just about playing the trend game. Over time, we're building tools from these trends, distilling them into useful building blocks that go beyond gimmicks and have a longer shelf life.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Phase 3: Innovate

Using AI internally within Lightricks

A/B testing

ZEEV: In the world of content creation, it's fascinating how AI has started to play a significant role. The actors you see in ads, and even in the setting or background, may not be real – they're created by AI. This allows for an incredible amount of A/B testing.

We're constantly measuring how people respond to creatives. In the marketing department, internal tooling that allows more A/B testing and asset creation is a huge thing. Every month, we allocate a certain performance marketing budget and task our team with achieving a Return on Ad Spend (ROAS) after 30 days, 60 days, and so on.

Using AI externally to innovate with customers

ZEEV: An example of how we applied AI innovation externally was in our partnership with a global entertainment company in the promotion and launch of their new fantasy heist action movie. This partnership was much more than adding a new feature to one of our apps. We wanted to transform how blockbuster movies were marketed and experienced by fans.

Objective

We created a campaign brief to foster a sense of inclusion in the movie's universe using the generative art technology in Lightricks's Photoleap app – so fans could create and share personalized character avatars in the likeness of one of the movie characters on social media. Lightricks already had a strong existing relationship with the global entertainment company and we aimed to jointly promote this feature, leveraging our partnership to offer fans a novel, AI-driven way to engage with the film.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

Technological challenge and creative solution

Designing AI avatars for the movie while retaining the user's unique identity was a particular challenge due to the multiple fantasy-based characters within the game's universe. We overcame this with a balanced approach, engineering design prompts and templates that captured the movie's essence while ensuring players could identify with their characters. The solution provided fans with an immersive experience that allowed them to see themselves in the context of the film's world without losing their personal touch.

Execution

During the premiere at SXSW, the entertainment company unveiled the Photoleap movie tool, engaging fans with an accessible avatar-creation feature via the Photoleap app. Fans could upload photos to craft avatars resembling their favorite characters, promoted through visual branding. Lightricks launched the campaign last year, boosting it with a paid social media campaign across multiple platforms and international markets to drive app downloads. Collaborative PR efforts between Lightricks and the entertainment company delivered widespread media coverage, aligning the campaign's technological innovation with the fans' passion for the fantasy action film universe.

Results and Impact

Using AI internally within Lightricks

ZEEV: By rethinking our core operating model and internal processes and making changes to serving key customer segments, we've made improvements throughout our business.

Our education of teams and how they access these new technologies has helped bring everyone on the journey together and embed changes.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

We've increased the speed and agility of our processes, especially internally. For example, our customer experience team now uses AI to respond to standard questions or issues, which is about 65% of the emails they receive. To give you some context, we have millions of monthly active users and tens of thousands of emails, so being able to reduce the volume to our customer support team by 65% really increases the personal attention to the remaining 35%.

AI innovation has driven significant marketing performance improvement, including internal tooling that allows more A/B testing and asset creation.

We're also doing more with fewer people and increasing our Return on Ad Spend (ROAS) targets as a big chunk of our performance marketing improvement was because of AI.

Impacts included:

- Increased speed and agility of internal processes, with 65% of customer emails now answered using AI.
- Improved performance marketing and increased Return on Ad Spend (ROAS) targets.

Using AI externally to innovate with customers

ZEEV: Because content creation is driven by trends in the consumer sector, we've tracked the trends and built new tools. These are useful building blocks that go beyond gimmicks and have a longer shelf life.

Our collaboration with the entertainment company was a huge success, demonstrating the powerful synergy between cinema and cutting-edge AI technology. Not only did it deepen fan engagement but it also demonstrated the vast potential of AI in creating personalized promotional strategies. The initiative reinforced Lightricks's position as a pioneer in providing large-scale, customized, and immersive experiences. It showcased our vision and helped cement our industry influence, and paved the way for future innovative campaigns.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

The campaign achieved substantial engagement, with 350,000-character bundles created and a 25% increase in app installs. By June 23, the campaign had generated significant media attention with 13 detailed stories in notable outlets, leading to over 15.6 billion potential impressions. At Lightricks we also harnessed social media engagement, reaching 20 million people and incorporating user feedback to iteratively refine the AI and improve the user experience.

Impacts included:

- Increased consumer engagement with a 25% increase in app installs.
- Enhanced brand reputation with 15.6 billion PR campaign impressions.

Key Learnings

- **Leadership understanding of AI technology:** Leadership with an intimate hands-on understanding of AI can crucially influence resource investment, project prioritization, and work approach. A proactive approach ensures optimal utilization of technology, keeping the business at the forefront of innovation.
- **Evolving business demarcations:** AI is eroding traditional boundaries between different businesses. Organizations should be prepared to adapt and evolve as these boundaries shift.
- **AI and code generation:** LLMs have matured to a point where they can generate high-quality code. The development of Multimodal Large Language Models (MLLMs) that can understand and manipulate visual elements like photos and videos has the potential to automate tasks that previously required human intervention. AI is a game-changer for software engineering.
- **Customer use of AI:** In the consumer sector, it's important to stay attuned to trends and develop tools accordingly. Trend-based tools or products can eventually become useful building blocks with a longer shelf life. At the same time, understand and stay close to how your customers are using AI in their workflows.

MULTIMODAL LLMS REDEFINE SOFTWARE DEVELOPMENT AND CUSTOMER INNOVATION

- **AI in internal processes:** Integrating AI into high-value internal processes can have a positive impact on the productivity of developers, marketing people, and designers. Maintain transparency when presenting new ways of working to clients to build trust.
- **Strategies for smooth integration:** Incorporating AI into business operations can be challenging, requiring careful planning, gradual implementation, and incentivized adoption. Ensuring access to the latest tools and identifying early adopters for AI projects will help avoid setbacks. Develop legal guidelines to clarify what tools are acceptable for different teams to use.
- **AI-driven A/B testing:** AI plays a significant role in content creation, allowing for extensive A/B testing. It can create alternative audience personas and creative assets, enabling marketers to test, and refine messaging for optimal performance. Internal tooling in the marketing department further streamlines the process by facilitating rapid A/B testing and asset creation.
- **Synergies between diverse expertise:** Multidisciplinary teams tend to outperform highly specialized teams, while interdisciplinary collaboration can open new avenues for innovation.

34

Where We've Got To and What's Next

AI is not a recent development, but extracting value from it remains a challenge for many businesses. The shift toward efficiency and improved productivity is a crucial aspect of modern business operations. This is where AI can play a significant role.

Learning how to leverage data and put AI to work has been the core purpose of this book.

Despite the wide acknowledgment of AI's benefits, many organizations and boards still struggle to use AI to generate value. Learning from leading AI practitioners has been key to meeting this challenge. *The AI Value Playbook* does not claim to have all the answers, but by sharing learnings on how AI can promote growth, improve cost efficiency, and mitigate risk, we can all benefit.

Over the past 10 years and, more recently, with the launch of publicly accessible LLMs, we have had clear evidence of what an AI-driven organization can achieve: making operations more resilient, improving core processes and customer experience, managing risk, and strengthening employee engagement.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Technological progress rarely arrives in a single, disruptive moment, but instead evolves through seemingly minor innovations over time. This often makes it difficult to foresee its full potential impact. The full impact of AI is yet to materialize primarily because large sectors of the economy are at different stages of integrating AI into their business practices. It will take time.

Some organizations are mature in their understanding and use of AI, others are just beginning their journey, and many are at various stages in between. Certain industries have higher switching costs and more significant regulatory hurdles. This broad range highlights that a flexible framework is essential to integrate AI into a business, one that can evolve and accommodate different levels of maturity and scale.

This final chapter is designed to encapsulate the essence of the entire book into a playbook in its own right – the AI Value Playbook (AIVP). Drawing on both my professional experience and the insights from interviews and case studies, I have developed the AIVP framework. This framework consolidates the successful approaches, common themes, and first principles identified across a diverse range of industries and business models, revealing underlying commonalities and learnings.

The AI Value Playbook

This playbook has been developed to guide organizations on how AI can be implemented to drive value. It's not about the "right" way; it's about finding effective solutions to complex problems and continuously learning.

The AIVP outlines four key interconnected components that need to be worked through to successfully implement AI:

- **Strategic Alignment and Target Application:** This ensures the AI initiatives align with the company's strategy. It prioritizes how AI will be used within the organization to measurably add value. This includes defining the AI use case – the specific problem it's designed to solve – and its interaction with other

WHERE WE’VE GOT TO AND WHAT’S NEXT

systems to achieve this goal. It also details the AI system’s inputs, outputs, and interactions, and measures the resulting benefits, which form a business case.

- **Technical Capability:** This starts with an internal evaluation of data quality, security, and technical capabilities. This assessment establishes a baseline and benchmarks key aspects of an organization’s readiness for AI implementation. It analyzes current resources, tools, practices, architecture, integration points, performance, and scalability. It serves as a foundation for developing a robust AI roadmap aligned to the business’s growth strategy and uncovers gaps, risks, and opportunities.
- **Operating Model:** The operating model serves as a blueprint for the effective delivery of an AI solution. It defines which resources are required, how they are organized and allocated, and the operational procedures that need to be followed.
- **AI Adoption and Managing Change:** This is the process of deploying and getting benefits from an AI solution. It’s about effectively integrating the AI solution into business operations and culture. It involves internal training, addressing resistance, continuous monitoring, and adjustment of the implementation process to ensure the solution is technically proficient, accepted, and used.

Strategic Alignment and Target Application	Technical Capability	Operating Model	AI Adoption and Managing Change
Strategic	Technical	Implementation	Adoption

Figure 34.1 – The four key components of the AIVP

STRATEGIC

Strategic Alignment and Target Application

Is the use of AI aligned with your business strategy? Will it solve a problem important to your organization?

Many businesses are grappling with the following situation: “We’ve got lots of data in a data lake, we want to extract value from it and use AI, but we’re unsure how to go about it.” The answer to this starts with strategic alignment, target application, and measurement.

Business strategy alignment and leadership play a vital role in successful AI delivery. It’s essential to have a clearly quantified AI strategy that evolves through an iterative process and addresses priority business problems. Alignment is crucial for realizing value, whether it’s about making or saving money or mitigating risk. This involves understanding how AI can continuously enhance and transform business operations, products, and services, improve customer experience, manage risk, and strengthen employee engagement.

Drawing from my firsthand experience, leadership has an important role to cut through competing priorities and spearhead AI as a driver of business growth. While working with a global infrastructure investment firm, where the talent is innately analytical and data literate, the CEO was instrumental in moving the AI program from “critical, do later” to “critical, do now.” He had a clear ambition to integrate AI into their operations, not only by modernizing the data estate and adopting new technologies but also by redesigning his business model to make the best use of these advancements. His ambition and vision brought everyone on board and made the difference.

Leadership is responsible for setting the direction, defining the strategy, objectives, metrics, and incentives. It fosters a culture of innovation and ensures that teams have the resources they need to execute AI initiatives. Success often comes through a cycle of trials and learnings. Leadership’s responsibilities include investing in talent development, infrastructure, operating model adjustments and research.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Moreover, leaders must also be able to articulate the value of AI to stakeholders, manage risks, and ensure the ethical and responsible use of AI.

Is the application of your AI project targeting the right need to make an impact?

Organizations rarely struggle to put forward a wishlist of problems they want AI to solve, but these need to be prioritized against capability and value impact. To create focus, bringing different stakeholders, including business users, technical teams, and finance colleagues into one room to brainstorm optimal use cases is enormously beneficial. This process involves structured problem-solving techniques and design methodologies to focus on the issues that really matter and where AI can generate value.

When determining the value of an AI solution, a key question to ask is: who owns it? While the technology team may build the solution, they can't generate value with it. Technical solutions are often developed hoping the business will find value; however, excluding the business perspective during development can lead to insufficient attention paid to user experience, adoption, and value measurement. The true value comes from the stakeholders being fully invested in making it work by using the solution and realizing its benefits. Ownership lies best with the business where the use case stems from.

Occasionally, we can get caught up in the pursuit of pushing the boundaries of AI capabilities and overlook the importance of clearly defining the problem we are attempting to solve. Absorption in the quest for solutions can overshadow the need for a strategic pause. It is important for the business's practitioners to conduct a root cause analysis and establish whether the problem can be framed and quantified. Understanding the organization involves examining the tasks and the processes the people follow. Collaborating with individuals who have a deep understanding of the domain, and closely observing their work, will avoid preconceptions. It's about asking the right questions to develop use cases (a specific application of AI to benefit the business), Objective and Key Results (OKRs), and capability requirements, including process and change impact and technology.

WHERE WE'VE GOT TO AND WHAT'S NEXT

It's essential that these use cases are easily understood by all stakeholders, that they provide a clear direction for the AI solutions and a measurable way to track.

People are more likely to use an AI solution if they see a clear personal benefit. When defining the use case and OKRs, it's crucial to be able to answer: "What's in it for them?" This answer should allow for iteration over time, enabling continuous value realization and improvement.

Once everyone agrees on the problem definition and the opportunity for AI – the use case, target design, and its potential value – you can then discuss the technical capabilities needed to support it and develop an execution plan.

How effective is ROI measurement? Is the investment worth the outcome?

While many organizations develop targeted AI use cases, few follow through by measuring their value and impact. It's vital to not just develop these applications, but to also keep an eye on their value through initially developing minimum viable product (MVPs), which focus on testing essential functionality and effectiveness. This is crucial for buy-in across the business and future scaling.

AI projects require an agile and outcome-focused approach; it's not just about adding more tech or more people. The focus shifts from a technical pipeline to delivering tangible business results.

A good approach is to take a financial P&L lens to solving a problem by assigning each team their own line item to answer the challenge. Adopting a "build-measure-learn" cycle is key to keeping this all aligned – building the solution, tracking its impact, and learning as you go for constant improvement. This loop ensures everyone is in tune and pursuing the right areas for success without chasing pipe dreams of perfection in AI models, which can lead to excessive costs and diminishing returns.

Demonstrating tangible monthly savings and incremental value will justify continued business investment. Success comes from starting small, securing executive buy-in, and guiding the business through AI adoption. This means

WHERE WE'VE GOT TO AND WHAT'S NEXT

being resource-savvy and starting with MVP models with a profit focus. This process requires time and patience.

The Individual Plays

Align leadership and strategy:

Articulate a clear business case by defining and quantifying the problem to solve and the gap between the current state and the proposed solution. Ensure the problem is manageable in scope for resolution. This will secure buy-in.

Develop a well-defined AI strategy ensuring leadership understands AI as a risk and value driver. Align this strategy with key organizational goals and targets, ensuring it encompasses the required infrastructure, resources, funding, and organizational support.

Ensure purposeful AI deployment by prioritizing projects that demonstratively contribute to achieving your organization's strategy and impact key value drivers. Maintain focus on business outcomes and the impact of every AI project.

Align incentives to the appropriate individuals with incentives for middle management aligned to successful AI implementation. Otherwise, they may resist the changes, which frequently cause failures in projects.

Shift ownership to beneficiaries, those in the business who stand to benefit most from an AI solution. This simple shift leads to a more efficient use of resources and a clear demonstration of value.

Maintain a clear vision and focus to keep project momentum, even amidst competing priorities.

Play the long game, making necessary investments in engineering, infrastructure, operating model changes, and partnerships. There can be quick wins along the way but the implementation of AI into core processes is not an overnight sensation.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Pinpoint value and prioritize:

Understand how AI disruption could impact your industry and existing business model. Analyze how AI could solve operational pain points and align with the organization's existing strengths. Leverage experienced AI practitioners and industry specialists to highlight the potential disruption from AI now and on the horizon.

Focus on the priority challenges the organization is facing to pinpoint where AI can enhance value. Which functions contain high-potential processes that could drive near-term savings? Where does data become the primary asset and a source of IP? What are the data risks? This involves examining people's tasks and processes to develop the best target use cases for commercial gain.

Work with both data and domain experts for the specific sector to ensure a clear framing of the commercial business opportunities/use cases for AI and the likely value impact.

Formulate and prioritize hypotheses that are practical and relevant to value creation. Establish a backlog of hypotheses depending on urgency, feasibility, and cost.

Develop a roadmap for implementation aligned with the organization's growth strategy and outline the practical steps for implementing your highest-priority hypotheses.

Develop a minimum viable product (MVP):

Ask the right questions to extract the most meaningful output, otherwise the dataset will just be another breakdown of results.

Focus on a manageable scope by selecting two to four hypotheses/business questions to focus the MVP on, depending on the size and number of data sources.

Build an MVP to validate priority hypotheses quickly. This outlines how data will be used, how it will be sourced and managed, and the potential AI solutions. This MVP can be developed within a typical 12-week delivery timeline. This includes bottom-up use case preparation, design and UI, build, UAT/QA, and deployment.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Plan for scalability, demonstrating incremental progress and value to stakeholders to secure ongoing investments and counter the impatience for impact. Think big, start small, scale fast.

Measure:

Shift to business-outcome-centric solutions by moving away from a technical focus to prioritize solutions that deliver clear business benefits. This involves considering objectives, key results, and methodologies that focus on continuous value realization and customer benefits.

Adopt a “build-measure-learn” cycle in an iterative process to help continuously improve the solution based on the learnings from each cycle.

Measure with a financial lens by taking a P&L approach to solving a problem. This helps the team see how their work directly impacts the business's bottom line.

Monitor, update, and maintain the developed solution to ensure the solution remains optimized and delivers long-term benefits.

TECHNICAL

Technical Capability

Does your organization have the capability to provide and manage the data foundations and services for AI?

The technology team leads the essential task of deploying AI by providing shared services to the data and AI teams in the business. The technology team can help the business optimize its operations and focus on revenue-generating areas. This starts with a clear evaluation of the current AI maturity, using a structure such as the following:

WHERE WE'VE GOT TO AND WHAT'S NEXT

Strategy

PEOPLE	AI DEVELOPMENT LIFECYCLE	TECHNOLOGY
Leadership owns a well-defined AI strategy aligned to key targets and includes the required infrastructure, resources, funding, and organizational support. Collaboration between the AI skills in the organization and other teams – AI is not siloed. Ability to translate business needs into AI technical requirements. Plan for training and development of the workforce in AI.	AI is used to create new value, enhance customer and employee interfaces, and optimize efficiency and productivity in operations. Employees have domain-specific knowledge relevant to the specific application of AI in the target use case.	Clear technical strategy, aligned with the business strategy, and a roadmap that defines the capabilities required to deliver the organization's AI strategy.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Delivery

PEOPLE	AI DEVELOPMENT LIFECYCLE	TECHNOLOGY
Proficiency in defining project scope and selecting appropriate algorithms. Competence in managing the development pipeline. Expertise in data collection, cleaning, transformation, and feature engineering. Ability to handle large datasets and ensure data quality. Skills in data engineering, machine learning, and software development and operations. Knowledge of model development strategies, monitoring, and integration into existing systems. Knowledge of relevant tools and frameworks for building and deploying AI models.	High-value business processes and decision-making are powered by AI. Proficiency in model training, hyperparameter tuning, and evaluation methodologies. Ability to evaluate model performance using relevant metrics and identify potential issues like overfitting.	Technology architecture and shared services enable the delivery of priority use cases. Data lake provisioning. Common data model and data catalog. Data access management. Source code and delivery processes are managed, shared, and checked for quality. Expertise in relevant machine learning libraries, cloud platforms, and hardware infrastructures.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Management and Governance

PEOPLE	AI DEVELOPMENT LIFECYCLE	TECHNOLOGY
Understanding of data security and privacy regulations. Understanding of legal and ethical considerations in data handling. Knowledge of data management principles and practices and AI governance strategies. Executive sponsorship for data management and AI governance initiatives. Knowledge of data versioning, bias mitigation techniques, and responsible AI practices. Proficiency in data quality management and data architecture design.	Data governance documentation with access requirements. Implementation of data governance processes to ensure data quality, compliance, and responsible AI practices. Monitoring of the model's performance in production and logging its activity for auditing purposes. Tracking of explainability, when feasible.	Data management tools, data pipelines, and security measures for handling sensitive data used in AI systems. Identification and mitigation of potential biases in the data and the model's output. Implementation of explainable AI techniques to make the model's decision-making process understandable. Ability to identify and address potential security vulnerabilities in AI systems.

Table 34.1 – AI maturity assessment criteria

This assessment needs to be adapted to the scale of an organization to identify opportunities, benchmark progress, and quantify the value of implementing AI. This includes evaluating data quality and its ethical use. The risk of not doing

WHERE WE'VE GOT TO AND WHAT'S NEXT

this leaves gaps in understanding the level of data maturity, quality, integrity, and associated risks, and where the organization needs to get to.

Early opportunities for impact

Often, the initial outcome of an assessment identifies opportunities for quick wins. Even in an organization burdened with legacy systems and a suboptimal architecture, an AI solution can work effectively with minimal integration. For example, an AI-powered recommender solution can be used in a customer service context to predict the next best action for customer care staff with minimal integration into a legacy CRM.

An important aspect of the assessment is to understand user workflows and the industry-standard technology that people rely on for critical processes. With due consideration, the business can integrate industry-standard solutions where appropriate.

The assessment can also reveal in-house gaps in capability, which can be rapidly augmented through forming strategic partnerships with specialist firms. Peer relationships with other businesses are another valuable source for sharing expertise from both within and outside the specific industry. These relationships can help address common challenges, drive innovation, and build on the progress made by others.

Building efficient and safe AI systems

While the leadership team is responsible for reducing time to business value, the technology team plays an instrumental role. They centrally manage standardized services, including creating and managing data interactions with applications, IoT, and other tools to create business-owned AI solutions.

A common practice is to establish shared services for AI use cases, including data landing, log management, governance rules, compliance, ethics, and security measures. The technology team ensures the data lake is well managed, curated, and cataloged, as well as enabling data merging and cleaning. This allows the

WHERE WE'VE GOT TO AND WHAT'S NEXT

data and AI team within the organization to create standardized AI solutions using configurable, reusable services without extensive coding.

A low-code or no-code approach is often used to create AI solutions more efficiently, employing a user-friendly drag-and-drop interface and pre-built templates. This system automatically applies data quality and master data service rules.

This approach assumes the infrastructure is cloud-based with governance built in, so solutions can scale up or down as needed, making value realization agile and efficient. Thus, compliance happens by design rather than enforcement. It's a complex shift that takes time to implement.

Creating standardized and configurable environments with automated tools and pipelines is also pivotal for AI iterative development. They allow the data and AI team to work in their preferred manner with a governed framework for tracking and measuring results. This environment can be easily automated using tools like YAML and DevOps pipelines, eliminating the need for manual intervention. It not only reduces costs but also establishes governance around the process, enabling better control and understanding of activities.

Testing throughout the entire lifecycle

To ensure successful and optimized AI implementation, prioritize rigorous testing throughout the entire lifecycle. This includes integrating human expertise (HITL) for quality assurance, especially during training data creation. A dedicated DataOps team streamlines testing and ensures high-quality data throughout the AI development process. To catch issues early and optimize performance before public release, leverage internal “super users” through a rigorous internal testing process (dogfooding).

Ethics and reputational risk

Ethical AI and reputational risk considerations are an emerging field, encompassing a broad range of practices and strategies. The integration of ethics into the AI solution lifecycle is a critical aspect of responsible AI development. There are a number of evolving frameworks, and a key thought leader is Professor Luciano

WHERE WE'VE GOT TO AND WHAT'S NEXT

Floridi, Oxford Internet Institute's (OII) Professor of Philosophy and Ethics of Information, and Director of the Digital Ethics Lab at the University of Oxford. In his recent book *The Ethics of Artificial Intelligence: Principles, Challenges, and Opportunities*, he identifies five core principles for ethical AI that should be fundamental considerations for business leaders:

1. **Beneficence:** AI should be designed and developed to benefit people by solving meaningful problems and improving efficiency.
2. **Non-maleficence:** AI should not harm people. This requires a thorough risk and impact assessment and involves building secure AI systems to prevent malicious use.
3. **Autonomy:** AI should respect human autonomy by augmenting human decision-making and aligning with privacy legislation.
4. **Justice:** AI should be fair and just and not perpetuate bias or discrimination.
5. **Explicability:** AI should be intelligible and accountable. This implies establishing mechanisms for accountability, such as audit trails.

The European AI Act, formally known as the Artificial Intelligence Act, is a significant piece of legislation currently being finalized by the European Union (EU). It aims to be the world's first comprehensive legal framework for regulating artificial intelligence (AI). This act will address three key areas:

- **Trustworthy AI:** Ensure AI respects rights, safety, and ethics (prevent bias, harm).
- **Risk-based regulation:** Stricter rules for high-risk AI (facial recognition, medical diagnosis).
- **Innovation and competitiveness:** Promote responsible AI development for a strong European AI market.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Horizon scanning for emerging technologies

To inject progressive thinking and keep pace with the speed of AI innovation, establishing a design authority is paramount. This forward-thinking team scans the horizon for emerging trends and technologies, ensuring competitive advantage. Furthermore, the organization needs to understand the value of building modularity into core capabilities. This future-proofs the organization by allowing components to be updated as the business needs evolve across different timeframes.

The Individual Plays

Assess current capability and quick wins:

- **Establish a baseline of data quality, security, and technical capabilities.** This involves evaluating architecture, integration points, performance, and scalability. This helps benchmark progress and identify opportunities for potential optimization and growth.
- **Plan risk mitigation and targeted investments** based on a thorough understanding of the foundational tech components: infrastructure, AI tools and algorithms, and security measures.
- **Define a 3-year tech roadmap** with a drill down on what needs to change in the next 6-9 months and set the tech capability journey in the right direction. Include both offensive and defensive investments.
- **Establish an ecosystem of AI specialist partners** to augment in-house capabilities and develop relationships with peers to collaborate on solutions to common problems and drive industry innovation.

Establish a robust data foundation:

- **Build data architecture capabilities** as the first step in setting up the foundation for AI solution development.
- **Establish data management and governance** practices to ensure high-quality, controlled data inputs for AI development. Integrating data and AI

WHERE WE'VE GOT TO AND WHAT'S NEXT

management allows the organization to concentrate on deriving value from their data in a controlled, compliant manner.

- **Implement data security and access controls** for sensitive data, restricting access to authorized personnel only.

Enable scalable and efficient development:

- **Ensure the infrastructure is cloud-based**, so solutions can scale up or down as needed, enhancing agility and efficiency with governance built in to ensure reliability and compliance. This approach optimizes value realization and drives business growth.
- **Enable shared services** for AI solution development. Build capabilities into the data architecture. Put in place mechanisms to source, cleanse, and control key data inputs and ensure data and AI management are integrated.
- **Implement architectural standardization** to enhance efficiency through automation and configurable, reusable capabilities.
- **Make regular updates and integrations** to enhance the model's capabilities, keeping it relevant and valuable. Integrate the AI solution with existing systems in the organization to provide a user-friendly experience for information access through familiar interfaces.

Consider solution design integrity:

- **Analyze user workflows** as well as their ingrained technology preferences, and build these into your solution, where appropriate, to increase engagement.
- **Ensure that AI technology decisions are both commercially viable and customer centric.** Conduct preliminary customer research to guide technology decisions, and then incorporate a continuous customer-centric approach with structured customer feedback systems.
- **Prioritize integrating AI into high-value processes**, while maintaining awareness of innovations that could impact other aspects of the value chain or decision-making, and swiftly adapt.

WHERE WE'VE GOT TO AND WHAT'S NEXT

- **Preserve technical modularity** in engineering and architecting. Carefully consider the “buy versus build” approach for each component, as the trend of AI service development will commoditize solutions.
- **Leverage domain specialist engineers** to work in collaboration with domain process designers, for successful proprietary model building.
- **Avoid over-engineering a model**, balancing mathematical rigor with practical human logic to solve the problem at hand. Understand the pros and cons of each method or algorithm. Plan your architecture, including data flow, before starting to code.
- **Take an iterative and user/customer-led approach** to the development of AI solutions. Integrating feedback and input is an ongoing process. Over time, this approach allows for ongoing enhancements and improvements.
- **Establish a design authority** that is responsible for scanning the horizon for emerging technologies in order to maintain a competitive advantage.

Evaluate the opportunities and impact of LLMs:

LLMs are emerging as a powerful tool, with significant potential to reshape various aspects of business operations for many industries. Some organizations are leveraging LLMs to solve long-standing problems, and other organizations are being left behind. The following are five key areas to explore the opportunities and impacts:

- **Human-Machine Interface (HMI)**, where LLMs serve as the front-end for interactions, offering a dynamic and prompt-led UI. This may lead to user prompting replacing static reporting, dashboarding, and fixed user interactions.
- **Programming**, where the rapid development of LLMs are impacting programming by aiding in learning to code, assisting in software engineer tasks, generating code, analyzing assisting in software engineering, analyzing large data sources, and enabling multilingual programming.

WHERE WE'VE GOT TO AND WHAT'S NEXT

- **Organization knowledge distribution and automation**, where LLMs are being used to collate enterprise knowledge, enhance employee productivity, and automate and standardize processes. This involves refining proprietary organization knowledge, enabling personalized and dynamic processes, and hyper-automation by no-code agents.
- **Data platforms**, where LLMs' ability to process and analyze large data sources could potentially expand the scope of data lakes. LLMs can handle diverse data formats and integrate information from various sources. LLMs can uncover hidden patterns and insights within the data, enabling organizations to make more informed decisions.
- **Broadening perspectives**, where LLMs can be trained on an organization's data and IP, so they can become a valuable tool to support the creative process. LLMs can generate a vast quantity of ideas, although the quality can vary. Therefore, human creativity remains crucial in the task of combining and transforming these ideas into innovative, high-quality solutions.

Prioritize testing throughout the lifecycle:

- **Integrate human oversight, Human-in-the-Loop (HITL)**, throughout the development cycle. Human expertise plays a crucial role in quality assurance, particularly when building training data for models. This ensures accuracy and robustness in your AI systems.
- **Implement a rigorous internal testing process, “dogfooding”**, for early quality checks. Leverage internal “super users” (knowledgeable employees) or select customers in a private preview to identify bugs and usability issues before public release.
- **Establish DataOps** for streamlined testing and quality assurance throughout the AI lifecycle. This includes implementing robust testing procedures to guarantee data quality and system performance.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Optimize performance post-launch:

- **Gather and analyze customer feedback** for continuous improvement through a combination of methods; select in-depth interviews with key customers, surveys, and analysis of customer support interactions.
- **Use a metrics-driven approach** to track model performance, customer engagement, Key Performance Indicators (KPIs), and overall solution efficacy.

Design trustworthy AI systems:

- **Monitor this rapidly evolving field** as considerations and legislation vary across geographies encompassing a broad range of practices and strategies.
- **Ensure AI solutions consider core principles for ethical AI**, encompassing regulatory compliance, system validity, safety, security, accountability, transparency, explainability, privacy, and fairness.

IMPLEMENTATION

Operating Model

How does your operating model enable the integration of AI solutions? Does it bring the right teams and skills together?

Often, when I work with business leaders, they underestimate the complexity of integrating AI and assume that hiring a data scientist will be a silver bullet that transforms their organization and future-proofs it for the next three years. However, driving value through AI requires diverse expertise. Multidisciplinary, cross-functional teams with software developers, data scientists, and data engineers are more likely to outperform highly specialized teams. They can handle a variety of use cases and build different types of AI solutions.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Capability needs will vary, but the composition of a high-level team includes:

- **Business architects:** Individuals who lead the overall approach and bridge the different domains of business and technology to define the strategy, governance, use cases, and alignment toward value realization.
- **Data engineers:** Specialists in moving, storing, and managing data within cloud environments. Responsible for delivering automated consistent datasets to AI systems.
- **Software engineers:** Coders who grasp software development principles.
- **Data scientists:** Analytical data experts who are closely tied to the use case. They focus on creating models to meet specific business needs and are, ideally, employed by the business unit they serve.
- **User experience professionals:** Includes both designers and UI experts who ensure the end product is functional and graphically appealing.
- **MLOps:** Machine learning operations engineers who understand ML algorithms and software deployment. They collaborate between data scientists and operations to manage the lifecycle of models.
- **DevOps:** Engineers who understand the software lifecycle, deployment, and maintenance of AI-enhanced software solutions.

Of these five capabilities, three are worth further focus. Firstly, the importance of data engineering teams cannot be overstated. They form the backbone of any data-driven operation, ensuring that the data is available at the right time and in the right format. These are critical to deploying numerous models.

Other crucial skill sets that are sometimes overlooked are business architects and User Experience (UX) professionals.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Business architects play a pivotal role in bridging the often significant gap between the technical and business teams. The technical team, while proficient in determining what needs to be built, may sometimes fall short in collaborating effectively with the business side to secure investment and ensure the adoption of the solutions. Conversely, business leaders may not always know the right questions to ask the technical team regarding the possibilities and means to achieve them. Business architects resolve this disconnect, integrating the technological possibilities with the business objectives.

UX professionals ensure the end product is functional, user-friendly, and engaging. Incorporating UX skills into your team can ease adoption and change management. UX professionals work at the beginning of the product design cycle and revisit their work with each improvement to the product. Investing in UX at the outset often leads to more successful outcomes. Additionally, incorporating design thinking into the team structure fosters a closer relationship between IT and business, resulting in products that are more useful to the business.

Supplementary infrastructure skills may not be required if cloud adoption is mature within the business, as many of the tasks traditionally performed by these roles can be automated. However, the type of AI solutions being built will influence the skill sets required within a team. For example, building complex AI systems may necessitate architectural skills to design robust and scalable solutions. In addition, governance capabilities are critical to ensure compliance with data regulations and ethical standards. Furthermore, the rapid evolution of large language model technology requires a blend of technical expertise and conceptual understanding in deep learning, NLP, prompt engineering, and architectural patterns.

Sourcing talent

Recruiting the right talent is challenging but investing time to get clarity on role requirements and a commitment to diversity can build teams that drive efficiency and value.

WHERE WE'VE GOT TO AND WHAT'S NEXT

In sourcing technical talent, it's important to consider domain expertise as well as the availability and cost of resources. Generally, resources for popular technologies like Python or Power BI are more abundant and therefore cheaper, while resources for niche products or smaller vendors may be scarce and more expensive.

Understanding these factors and individual motivations can help you recruit and retain the right talent. Recognizing that not all individuals are the same – some may be starters, others may be finishers or maintainers – helps form a well-rounded team, while balancing skills across project phases is essential.

Data scientists have very different backgrounds, from Oxbridge PhDs to data engineers with experience in using standard libraries. The task requested of the data scientist needs to align with their individual skill set and experience. If a high-flyer is given routine tasks, they are likely to feel a lack of learning and growth, stagnate, and leave. Equally, giving a complex and ambiguous task to a less experienced data scientist may leave them frustrated and unmotivated.

Talent shortage is a commonly cited problem. However, there are large pools of passionate and motivated professionals in all forms of technology. Seasoned tech professionals I speak with are eager to learn, “How can I stay up-to-date with AI innovation?” The real challenge is in hiring, training, and retention. For example, there's no value in asking for 10 years of experience in a technology or tool that hasn't been on the market that long. The real questions are, do employers really understand what is involved in the work? How to find and develop these skills? How to upskill the existing workforce?

Traditional methods, based on CVs from familiar talent pools, are insufficient due to skill shortages and competitive rates of compensation. It's important to seek motivated technical developers to join the team, focusing on finding passionate self-starters. Using practical assessments like hands-on building and reviewing models can help gauge immediate technical skills, while asking for feature ideas helps identify those who are really vested in the development of your organization. To nurture and retain talent, pair junior hires with senior team members to help them up-skill into professional developers. Implement structured processes like design documents and retrospectives to promote continuous learning.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Academia is a popular recruitment channel, offering a diverse range of talent to various industries. However, one challenge that often arises is the inconsistency in the use of tools and techniques. Academia exposes students to multiple open-source tools and methodologies, each with its own features and applications. While this process broadens their skill set, it can also lead to a lack of uniformity when they transition into the workforce. A potential solution is to create contained environments for data scientists to work with a variety of tools and techniques, while also promoting the development of common libraries and reusable components for a more consistent approach to problem-solving and the development of AI solutions across the workforce. It's a win-win situation that enhances both individual growth and collective productivity.

Organizing talent

Many organizations begin with a distributed model for AI skills, which offers speed, flexibility, and the ability to tailor services and solutions. However, this approach can lead to several drawbacks, including varying maturity levels, redundant efforts and investments, missed opportunities for the cross-fertilization of practices, and higher overall costs.

At the other extreme, a fully centralized model offers economies of scale, common governance and risk management, and shared resources but can be slow to respond to specific business needs.

The proposed hub-and-spoke model provides a more efficient and adaptable approach, aligning well with evolving business models. It leverages the benefits of central oversight, governance, and risk management while enabling the consolidation of investments and promoting a degree of autonomy and agility within the business. This model features a lean Center of Excellence (COE) with expertise in data science, data engineering, operations, governance, and data leadership. The COE sets standards for solution development, ensures repeatability, and promotes advanced capabilities. Hands-on AI teams, often including data scientists and analysts, are integrated into the business, to ensure agility and alignment with specific needs.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Board-level AI expertise

Effectively integrating AI necessitates expertise at all levels of an organization, including the board. While some board members may possess this expertise, others may need to develop it. Building this expertise could come from recruitment, training, and engagement with external advisors. The value this expertise can bring includes strategic guidance on data management and AI investments, insights into technology choices, and implications of emerging innovations. It can help an organization stay ahead of the curve and make informed decisions, thus accelerating its path to becoming a modern, AI-driven enterprise. In addition, it's crucial for boards to consider the ethical implications of AI, such as data security and privacy, algorithmic bias, and the impact on jobs.

The Individual Plays

Establish a team with diverse skills:

- **Prioritize domain expertise** as a key differentiator.
- **Deploy multidisciplinary teams** as these tend to outperform highly specialized teams.
- **Consider a hub-and-spoke operating model** with a lean Center of Excellence (COE) containing AI specialists including data scientists, engineers, operations, governance, and analysts. The COE provides expertise and support, while embedding AI talent within business units promotes ownership and trust and ensures responsibility for delivering value.
- **Embed AI technical skills** including data scientists, engineers, and analysts into business units to foster ownership and trust and ensure responsibility for delivering value.
- **Hire specialist engineers with domain and tech expertise** to work with domain process designers, for successful proprietary model building.

WHERE WE'VE GOT TO AND WHAT'S NEXT

- **Ensure engineering talent is prioritized** to meet the growing demand for engineers skilled in integrating complex data systems. Given the rarity of full stack data scientists, borrowing engineering talent from other teams can deliver the right blend.
- **Hire adept bridge people (business architects) who** can understand and communicate across different domains, to facilitate cooperation and mutual understanding among various teams.
- **Integrate UX expertise** to drive adoption and add value to AI projects.
- **Prepare for a hybrid workforce** where AI augments workflows. Success follows when leaders embrace an operating model that leverages the strengths of both humans and machines, rooted in agility, flexibility, and continuous learning.

Prioritize diverse sourcing strategies and practical assessments:

- **Diversify hiring strategies** to consider candidates from various backgrounds and educational institutions. Traditional methods are insufficient due to skill shortages and competitive rates of compensation.
- **Seek out motivated technical developers** and focus on finding passionate self-starters.
- **Emphasize practical assessments** such as hands-on model building or reviews to assess immediate technical skills. This approach fills skill gaps but also promotes diversity and innovation.
- **Nurture and retain talent** by pairing junior hires with senior team members to help them up-skill into professional developers. Additionally, implement structured processes like design documents, code reviews, and retrospectives to promote continuous learning.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Build board-level expertise:

- **Ensure AI expertise within your board** through recruitment, training, or engaging external advisors. An AI-conversant board can impact the organization's strategy, operations, and competitive landscape, as well as considering the ethical implications of AI.

ADOPTION

AI Adoption and Managing Change

What strategies and processes does your organization have in place to manage the transition toward AI adoption? Are your employees equipped and ready?

In a recent meeting with a senior healthcare leader, I was asked to discuss a failure in an AI project. The example that came immediately to mind was a failure in adoption and change management in a company where it was assumed everyone was onboard with the objective to become an AI-led organization – and committed to the work needed to get there. However, the reality proved otherwise. This failure in adoption and change management is a common issue, echoed repeatedly throughout this book and by other practitioners. Successful AI implementation has both a technical aspect and a human element. The involvement of people within the organization as a linchpin of successful AI adoption is repeatedly evidenced in our interviews.

A *Harvard Business Review* article by Randy Bean (packt.link/hbr) revealed the struggle of many Fortune 1000 companies. Bean highlighted that the barrier to becoming data-driven is not just legacy tech but embedded cultures that are resistant to new ways of doing things. Interviews and case studies in this book support this notion, revealing a stronger resistance in longer-established organizations. To overcome this challenge, Bean suggests a solution lies in engaging the people who will work with and be affected by AI in every step of the process.

WHERE WE'VE GOT TO AND WHAT'S NEXT

Focusing on the people and processes often yields quicker and more successful results. Conversely, projects that lead with technology tend to be slower and have a lower success rate. Driving adoption and ownership of the solution or model within the organization is crucial – it's not a case of “build it and hope they'll use it.” Users need to have a vested interest in the creation of these solutions because having “skin in the game” helps them realize its value.

The successful application of AI requires clear executive directives in a top-down approach and a comprehensive change management program. Active participation from every level of the organization in a bottom-up approach is also essential. Understanding how to implement these changes in the details of daily work is also needed. Questions such as “Why is this important to my job?”, “What do I need to do differently?”, and “Can I trust it?” should be addressed. It's crucial to ensure a team understands the benefits of the new tools or solutions, how to use them, and the improvements they can bring to their work. Engaging with the people who will ask these questions is critical, making them active agents of change rather than passive recipients.

To overcome resistance, frame the project as an experiment rather than a fixed end goal and understand the limitations of your AI solutions. Respect the fact that the learning process is continuous and accept it includes failures that lead to incremental successes.

The business must own the solution, with a clear and concise use case that allows for incremental improvements and measurable progress. If successful, there's a virtuous cycle as the business starts saving money, stakeholders become satisfied, and they begin to use more AI solutions and capabilities.

Adoption and change begin with early initiatives aimed at achieving quick wins in areas where the value of the AI solution can be immediately demonstrated. Once these initial successes are achieved, the changes can be scaled throughout the rest of the organization. Success often starts with enhancing efficiencies in operational areas like finance, HR, and customer service. Learnings are subsequently extended to market products and services. This approach ensures a systematic and effective adoption of change. Success in one area often sparks change in others.

WHERE WE'VE GOT TO AND WHAT'S NEXT

The Individual Plays

Secure early wins and drive user adoption:

- **Start in functional areas** such as finance, HR operations, or customer service, focusing on cost savings and efficiency gains to provide extra funds to invest in AI project changes and hiring new people.
- **Involve users from the start** to drive adoption and ownership within the business domain, making sure there's a shared understanding of the benefits and a vested interest in its success.
- **Frame AI projects as iterative experiments** to secure buy-in. Ensure initial results align with key metrics and impact business goals. Understand the limitations of the data/tools and respect that the learning process is a continual one, with failures often leading to incremental successes.

Manage change through effective communication and training:

- **Harness storytelling and alignment** to create compelling narratives, with the customer in the role of the hero. Align data projects with the company's north star metric and identify leading indicators for desired outcomes.
- **Create an engaging environment for communication** that captures attention, encourages discussion, and responds to concerns. Address the last stage of the adoption of an AI solution first, for a clearer development and deployment approach.
- **Build trust** and guide the organization with patience, humility, and strategic quick wins. Sharing AI solutions in development with stakeholders – the 80/20 rule – is important to gain buy-in and accountability. Launch with a partial product to allow for improvements as interest grows.
- **Invest in continual learning and upskilling** to prepare for the AI future. David Shrier, Professor of Practice, AI & Innovation at Imperial College Business School, goes beyond technical skills and proposes five key principles to cultivate greater cognitive flexibility in your organization, including consistent practice, reflective thinking, gradual progress, peer learning, and creative exploration.

In closing

This book has explored the current AI landscape, driven by a strong interest in the potential of AI to propel sustained productivity growth across industries. We opened with the thesis that sustained productivity growth is, in part, contingent on the effective integration and board adoption of AI solutions across a range of sectors.

Tech leaders such as Amazon, Google, Netflix, Meta, Microsoft, and OpenAI have been investing in AI for over a decade, and their products have become integral to our daily lives. However, AI hasn't yet been put to work by wider industries and translated into efficient growth.

We are still in the early stages of implementing AI in business operations. Large enterprises with established partnerships with the tech leaders, gaining access to cutting-edge AI tools and expertise, have a clear advantage. For some businesses, particularly newer companies built on cloud-based technology, AI adoption comes naturally. However, for others, creating value and gaining a competitive edge through AI might be more challenging. Technology and investment alone aren't enough to drive transformation. We are likely to see a significant impact over the next 5 to 10 years. Therefore, the crucial question remains: how can organizations best integrate AI to achieve significant productivity gains?

AI is a rapidly evolving domain. Its transformative potential in business, the need for its ethical application, and the innovative leadership required all present significant challenges for organizations. These challenges are particularly pronounced when deciding how best to invest in AI technologies. Yet, as we've seen throughout the interviews, technology innovation and monetary investment alone are not sufficient to drive transformation. Integrating AI requires a long-term perspective. It takes sustained effort and ongoing iterations. Failures are inevitable, yet they can pave the way for incremental successes and continuous improvement.

WHERE WE'VE GOT TO AND WHAT'S NEXT

While financial results might not immediately reflect a lack of technology modernization and AI integration, the consequences are likely to become more apparent over time as organizations struggle to scale, keep pace with the market, and adapt to changing customer behaviors. Therefore, investing in technology can be evaluated as a forward premium, a payment into the organization's future valuation. Newly emerging helpful agents are poised to become a leading driver of AI adoption. GPT-4 presents exciting possibilities for the growth of virtual agents with more natural interactions due to advanced reasoning and response capabilities, and improved task handling through better accuracy and custom code generation.

If you came to this book uncertain or hesitant about how to put AI to work in your organization, I hope your doubts have been dispelled. Informed by my experiences and those of AI leaders and CEOs, you now possess a distilled set of core principles to guide you: the AI Value Playbook. I trust that the insights revealed will enable your organization to target and prioritize where AI can create the most value, making sure you have the right technology, talent, and culture in place with responsible AI governance. As your organization evolves, the details may change but the core principles will remain consistent and help inspire your organization to map a route forward.