

# Generative AI for Leaders Playbook

Applied Approach from AI Strategy  
to Responsible Innovation in the Enterprise



Adrián González Sánchez

# Generative AI for Leaders Playbook

Applied Approach from AI Strategy to Responsible  
Innovation in the Enterprise

**Adrián González Sánchez**



# Generative AI for Leaders Playbook

Copyright © 2026 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.

This book was written by Adrián González Sánchez. Generative AI tools were used only to assist with ideation, phrasing, and diagram drafts, and all technical content and code were created, verified, and tested by the author and Packt's editorial team. Packt does not accept AI-generated content that replaces expert authorship.

**Portfolio Director:** Pavan Ramachandani

**Relationship Lead:** Ali Abidi

**Project Manager:** Divij Kotian

**Content Engineer:** Aditi Chatterjee

**Technical Editor:** Simanta Rajbangshi

**Indexer:** Rekha Nair

**Production Designer:** Salma Patel

**Growth Lead:** Prajwal S

First published: August 2026

Production reference: 1200826

Published by Packt Publishing Ltd.

Grosvenor House

11 St Paul's Square

Birmingham

B3 1RB, UK.

ISBN 978-1-80638-849-3

[www.packtpub.com](http://www.packtpub.com)

*Dedicated to my wife, Malini, and to her mother, Sukon Runtasewee, now at peace. May it serve both as a tribute and as an expression of gratitude for showing me what it means to fight and persevere.*

*Adrián*



# Contributors

## About the author

**Adrián González Sánchez** is a professional with... no, wait. Let me explain my story in the first person.

At this point in my career, I consider myself lucky to have worked for and with international organizations across Europe, LATAM, and North America. My engineering background allowed me to start my career in technical and consulting roles, but I progressively moved into what I call hybrid, technical management responsibilities. I have done so across all kinds of data science, AI, and cloud computing firms.

I'm not the most experienced professional out there, but I have lived multiple professional lives in parallel by building my career in data and AI companies while still making time to teach at the university and executive education levels in Canada and Spain, manage master's and executive programs, write multiple books with international publishers such as Packt, O'Reilly Media, and ANAYA Multimedia, coach companies and executives around the world, and mentor thousands of international leaders through the GetSmarter for MIT Sloan and Harvard VPAL since 2018, including the Artificial Intelligence: Implications for Business Strategy program. And let's not deny that the professional reputation I have built over the last 15+ years has been exponentially strengthened by my online thought leadership and the visibility of my AI roles at Microsoft.

That curiosity and ability to multitask have exposed me to a myriad of projects, organizations, and both successful and failed AI initiatives. Every few years, I try to write down and consolidate the learnings and methodologies that I believe may be useful to others.

As a result, and as my seniority has evolved over time, I started writing more technical books focused on the specifics of AI management at the tactical project and product levels, eventually expanding my focus to executive-level best practices and approaches to AI management. After so many discussions and experiences, I felt it would be valuable to bring you something different, tangible, and adapted to the complexity of leading a company in the era of generative AI. Because I work across all three levels (strategic, tactical, and technical), I try to connect everything to provide you with a practical, realistic, top-down approach that helps you scale your AI initiatives much faster than would otherwise be possible. All without the "yada yada" of the current times, without selling anything, and without falling into clichés or superficial insights.

Let's see whether I deliver on that promise, but I truly believe you will enjoy this book. More importantly, I hope you will use it and draw valuable lessons from my experiences, chapter by chapter, so you feel that you are not taking a leap of faith, but following the recommendations of someone with real, practical experience in the AI industry.

*Few people know the true story behind this book. Beyond my personal motivation to share my tips, experiences, and recommendations, the reality is that I wrote the manuscript during one of the most difficult periods of my life. More specifically, much of it was written across several wards of Siriraj Piyamaharajkarun Hospital in Bangkok and at the palliative care center where my wife's mother spent the final months of her life.*

*It may seem contradictory that, during such a personal and difficult time, my mind sought an escape from the long, cold days spent in hospitals and somehow allowed me to write these pages, with greater or lesser skill.*

*In truth, I believe it was possible because, by comparison, everything happening around me seemed far more complex and difficult to endure. And so, chapter by chapter, I managed to bring all my ideas together and turn them into the work you now hold in your hands.*

*For this reason, and for so much more, this book is dedicated to my wife, Malini, and to her mother, Sukon Runtasewee, now at peace. May it serve both as a tribute and as an expression of gratitude for showing me what it means to fight and persevere, even during the most difficult moments.*

*And, of course, my thanks to my parents and my sister for their unconditional support, and to Ali Abidi and the outstanding team at Packt for placing their trust in me once again.*

*Adrián*

## About the reviewer

**Vasu Raj Jain** is a Lead Software Engineer at Amazon Ads with over 10 years of experience in large-scale distributed systems and advertising technology. He builds AI-powered ad-serving systems that generate billions of dollars in revenue and deliver ads at massive scale, serving live events with over 20 million concurrent viewers. His work spans programmatic guaranteed deals, share of voice, real-time bidding, and the core platform infrastructure that supports them. He contributes to IAB Tech Lab working groups that shape industry standards, mentors engineers, and writes about distributed systems, ad tech, and AI. A University of Florida alumnus, he previously built computer vision tools at Prime Video.

# Table of Contents

|                |             |
|----------------|-------------|
| <b>Preface</b> | <b>xiii</b> |
|----------------|-------------|

|                                   |     |
|-----------------------------------|-----|
| Free benefits with your book..... | xxi |
|-----------------------------------|-----|

|                                                       |          |
|-------------------------------------------------------|----------|
| <b>Part 1: Building Your AI Leadership Foundation</b> | <b>1</b> |
|-------------------------------------------------------|----------|

|                                                                               |          |
|-------------------------------------------------------------------------------|----------|
| <b>Chapter 1: Embracing Top-Down AI Approaches: The Leadership Imperative</b> | <b>3</b> |
|-------------------------------------------------------------------------------|----------|

|                                                             |   |
|-------------------------------------------------------------|---|
| Understanding what artificial intelligence means today..... | 4 |
|-------------------------------------------------------------|---|

|                                                                      |   |
|----------------------------------------------------------------------|---|
| Assessing generative AI for innovation and executive engagement..... | 6 |
|----------------------------------------------------------------------|---|

Creating your first AI Policy • 7

Defining your new AI committee • 17

Confirming your North Star for AI use cases • 25

Using Porter's Generic Strategies Framework • 25

|                                                |    |
|------------------------------------------------|----|
| Understanding the levels of AI management..... | 30 |
|------------------------------------------------|----|

Improving continuously via management systems • 32

Evaluating the AI management levels • 35

|               |    |
|---------------|----|
| Summary ..... | 42 |
|---------------|----|

|                                                                             |           |
|-----------------------------------------------------------------------------|-----------|
| <b>Chapter 2: From Turing to ChatGPT: A Brief History and Context of AI</b> | <b>43</b> |
|-----------------------------------------------------------------------------|-----------|

|                             |    |
|-----------------------------|----|
| Historical perspective..... | 44 |
|-----------------------------|----|

Pre-1950: Imagining intelligent machines before AI existed • 45

The 1950s: The birth of AI as a discipline • 47

The 1950s to the early 1970s: Optimism, funding, and early limitations • 49

Up to 1980: The first AI winter • 50

The 1980s: Expert systems, revival, and the second winter • 52

Between the 1990s and 2006: From hype to real-world impact • 53

Between 2006 and the 2010s: The deep learning revolution • 55

The 2020s: AI goes mainstream • 57

Reading the pattern: What AI history teaches your business • 59

|                                                     |    |
|-----------------------------------------------------|----|
| Key milestones from an AI adoption perspective..... | 60 |
| How generative AI changes your game.....            | 60 |
| Comparison with the printing press .....            | 64 |
| Lessons from past hype cycles.....                  | 66 |
| Geopolitics and power in the context of AI.....     | 68 |
| Key takeaways • 76                                  |    |
| Summary .....                                       | 76 |

---

## **Chapter 3: Exploring Key AI Terminology and Concepts** **77**

---

|                                            |    |
|--------------------------------------------|----|
| Understanding the levels of AI .....       | 78 |
| Narrow AI (NAI) • 78                       |    |
| Artificial General Intelligence (AGI) • 79 |    |
| Artificial Superintelligence (ASI) • 80    |    |
| The current status • 80                    |    |
| Main AI capabilities .....                 | 81 |
| Forecasting • 82                           |    |
| Classification • 82                        |    |
| Detection • 82                             |    |
| Completion • 83                            |    |
| Content creation • 83                      |    |
| Automation • 84                            |    |
| Optimization • 84                          |    |
| AI learning .....                          | 87 |
| Types of learning patterns • 88            |    |
| Data splitting • 89                        |    |
| The AI lifecycle .....                     | 90 |
| Technical operations.....                  | 92 |
| Deep dive into ML models.....              | 95 |
| Unsupervised ML models • 95                |    |
| <i>Anomaly detection</i> • 95              |    |
| <i>Clustering</i> • 96                     |    |
| Supervised ML models • 96                  |    |
| <i>Regression</i> • 96                     |    |
| <i>Classification (binary)</i> • 98        |    |
| <i>Classification (multiclass)</i> • 99    |    |

|                                                                           |            |
|---------------------------------------------------------------------------|------------|
| Deep dive into DL models.....                                             | 102        |
| Key considerations • 103                                                  |            |
| Types of DL models • 104                                                  |            |
| Illustrative use cases .....                                              | 106        |
| Summary .....                                                             | 113        |
| <b>Part 2: From AI Understanding to Business Transformation</b>           | <b>115</b> |
| <b>Chapter 4: Understanding the New Era: Generative AI and Agentic AI</b> | <b>117</b> |
| What is new in generative AI? .....                                       | 118        |
| Types of generative AI capabilities • 119                                 |            |
| What is Agentic AI? • 121                                                 |            |
| Key building blocks • 123                                                 |            |
| Typical adoption and customization patterns • 126                         |            |
| From traditional NLP to generative AI.....                                | 132        |
| Understanding key AI concepts • 132                                       |            |
| AI model training phases • 134                                            |            |
| <i>Pre-training</i> • 135                                                 |            |
| <i>Post-training</i> • 139                                                |            |
| <i>Inference</i> • 142                                                    |            |
| <i>Customization</i> • 142                                                |            |
| <i>Optimization</i> • 143                                                 |            |
| Identifying and prioritizing AI use cases .....                           | 146        |
| Use case discovery • 146                                                  |            |
| Use case evaluation • 148                                                 |            |
| Prioritization models • 149                                               |            |
| Summary .....                                                             | 151        |
| <b>Chapter 5: Crafting an AI Strategy: From Vision to Execution</b>       | <b>153</b> |
| The pillars of an AI strategy.....                                        | 154        |
| Assessing your organization's AI maturity .....                           | 158        |
| Sharing your strategic AI roadmap .....                                   | 162        |
| Turning your AI strategy into operational actions.....                    | 165        |
| Approaching the next chapters optimally .....                             | 166        |

|                                                                                |            |
|--------------------------------------------------------------------------------|------------|
| Summary .....                                                                  | 172        |
| <b>Chapter 6: AI-First Approaches: New Business Models and Products</b>        | <b>173</b> |
| Understanding the levels of AI transformation .....                            | 174        |
| Rethinking products and services .....                                         | 177        |
| The impact on product development • 177                                        |            |
| Implementing AI-enabled services • 178                                         |            |
| Becoming an AI-first organization • 179                                        |            |
| The pursuit of scalability.....                                                | 182        |
| Organizational scalability • 182                                               |            |
| Technical scalability • 184                                                    |            |
| The AI washing trap .....                                                      | 186        |
| Managing AI cost units and pricing models .....                                | 188        |
| The AI cost structure • 188                                                    |            |
| The token economy in the era of generative AI • 191                            |            |
| Summary .....                                                                  | 194        |
| <b>Part 3: Building the Organization for AI at Scale</b>                       | <b>195</b> |
| <b>Chapter 7: Building and Scaling AI Teams: Roles, Reskilling, and Hiring</b> | <b>197</b> |
| Your own role as an AI leader .....                                            | 198        |
| Own the AI narrative • 198                                                     |            |
| Establish visible AI thought leadership • 199                                  |            |
| Design pragmatic AI interview processes • 199                                  |            |
| Enable feedback loops for the AI initiative • 200                              |            |
| Tell a credible long-term story • 200                                          |            |
| Analyze the market and understand the investment • 201                         |            |
| Empower your AI teams • 202                                                    |            |
| Understanding the evolution of roles in the AI development cycle.....          | 202        |
| Traditional roles • 203                                                        |            |
| Frontend Software Developer • 203                                              |            |
| Backend Software Developer • 204                                               |            |
| Cloud Administrator and Engineer • 204                                         |            |
| Data Engineer and Architect • 204                                              |            |

---

|                                                                                    |                |
|------------------------------------------------------------------------------------|----------------|
| <i>Data Scientist and ML Engineer • 205</i>                                        |                |
| <i>Quality Assurance and Security Engineer • 205</i>                               |                |
| <i>UX Designer • 205</i>                                                           |                |
| <i>Product Analyst/Data Analyst • 205</i>                                          |                |
| <i>Other tactical roles • 206</i>                                                  |                |
| New roles • 206                                                                    |                |
| <i>Generative AI Developer • 206</i>                                               |                |
| <i>Generative AI Cloud Expert • 206</i>                                            |                |
| <i>Generative AI Architecture • 207</i>                                            |                |
| <i>Prompt Engineering • 207</i>                                                    |                |
| <i>Generative AI PM and others • 207</i>                                           |                |
| <i>AI Regulation and Responsible AI Expert • 208</i>                               |                |
| <i>Forward Deployed Engineer • 208</i>                                             |                |
| <i>Other critical roles • 209</i>                                                  |                |
| <b>Other considerations for your AI talent strategy .....</b>                      | <b>209</b>     |
| Shifting from science to engineering • 209                                         |                |
| Selecting between fundamental and applied research • 210                           |                |
| Avoiding trends and focusing on skills and responsibilities • 210                  |                |
| <b>Key considerations to scale your AI teams .....</b>                             | <b>212</b>     |
| Analyzing talent gaps and planning • 212                                           |                |
| Reskilling the current workforce • 213                                             |                |
| Addressing the future of work issues • 213                                         |                |
| Ignoring the noise and avoiding the AI layoff trap • 214                           |                |
| <b>Defining the operating model .....</b>                                          | <b>214</b>     |
| <b>Summary .....</b>                                                               | <b>218</b>     |
| <br><b>Chapter 8: Navigating the AI Ecosystem: Tools and Vendor<br/>Evaluation</b> | <br><b>219</b> |

---

|                                                    |            |
|----------------------------------------------------|------------|
| <b>Understanding the AI technology stack .....</b> | <b>220</b> |
| AI infrastructure • 222                            |            |
| Platforms • 222                                    |            |
| Development studios • 223                          |            |
| Libraries and frameworks • 223                     |            |
| Governance, compliance, and safety • 223           |            |
| <b>Evaluating tools and vendors .....</b>          | <b>224</b> |



---

|                                                                                             |                |
|---------------------------------------------------------------------------------------------|----------------|
| Defining the AI tooling selection criteria • 224                                            |                |
| <b>Leveraging cloud technologies for AI strategy .....</b>                                  | <b>234</b>     |
| <b>Building resilience while avoiding vendor lock-in .....</b>                              | <b>238</b>     |
| Technical flexibility and governance • 239                                                  |                |
| Negotiation, exit strategies, and open-source alternatives • 240                            |                |
| <b>Summary .....</b>                                                                        | <b>242</b>     |
| <br><b>Chapter 9: Responsible AI Implementation: Ethics, Governance, and Accountability</b> | <br><b>243</b> |

---

|                                                       |                |
|-------------------------------------------------------|----------------|
| <b>From AI ethics to accountability.....</b>          | <b>244</b>     |
| Key terms and principles • 245                        |                |
| <b>Quantifying RAI initiatives.....</b>               | <b>248</b>     |
| <b>Relation between ethics and regulations.....</b>   | <b>252</b>     |
| Types of regulations • 253                            |                |
| The rising role of AI standards • 261                 |                |
| <b>Notion of data and AI governance .....</b>         | <b>263</b>     |
| Role of the RAI Champions • 266                       |                |
| Illustrative toolkit to implement RAI • 267           |                |
| <b>Turning insights into action.....</b>              | <b>269</b>     |
| <b>Summary .....</b>                                  | <b>272</b>     |
| <br><b>Chapter 10: Unlock Your Exclusive Benefits</b> | <br><b>275</b> |

---

|                                                       |     |
|-------------------------------------------------------|-----|
| Unlock this Book's Free Benefits in 3 Easy Steps..... | 276 |
|-------------------------------------------------------|-----|

---

|                                  |            |
|----------------------------------|------------|
| <b>Other Books You May Enjoy</b> | <b>280</b> |
|----------------------------------|------------|

---

|              |            |
|--------------|------------|
| <b>Index</b> | <b>283</b> |
|--------------|------------|

---

# Preface

*Time for AI Leadership is now.*

*Adrián*

Saturday morning. It is a rainy day, and I'm sitting on a plane, flying and getting ready to spend one week at one of the main technology events of the year in San Francisco. During the last few weeks, I have received so many requests to meet and discuss generative AI adoption topics at both organizational and technical levels, from an interesting mix of big corporations, growing startups, and even public sector entities. All of them are led by amazing folks who are driving the unstoppable pace of innovation in their companies. You may be one of them. Who knows, maybe you are right now flying to this or another international event, hoping not only to develop interesting connections, but also to learn techniques you can leverage in your day-to-day activities.

While we are here, let me ask you a few questions. Are you truly leading your technology and AI adoption initiatives? Or are you just sponsoring and supporting them, then delegating or even outsourcing not only the what, but also the how and when? Can you go deep enough, as some sort of full-stack AI leader who can connect board-level discussions with technical implementations? Can you, executive friend, proactively guide your teams to follow a structured approach that still allows them to be the main character of the movie, but gives you certainty and clarity and reduces the probability of failure, one step at a time?

Continuing with this mix of interrogation and monologue, what if, instead of relying on traditional bottom-up, technical approaches to AI adoption (and technology adoption in general), technology leaders and C-suite individuals like you could not only lead but also define the approach and scope of how your company makes sense of the current wave of generative AI technologies? What if your role could become something relevant beyond that of a regular executive sponsor with high-level involvement?




Figure 0.1 – My very initial notes and sketches to ideate CEO/leader-led approaches to AI

Even more, imagine you could not only upskill to catch up with the core aspects of AI technology, but also proactively bring a structured approach that connects all the required dots within your organization and makes sure that any AI investment is worthwhile and carefully handled to minimize the inevitable risks related to an evolving set of technologies and the continuous innovation dilemma.

These questions have been on my mind for a while, like that favorite song that I play again and again, since I started working on data and AI projects more than a decade ago. They have also come up during my recent years of mentoring and tutoring companies and executives looking for a proper AI strategy, during which I have witnessed the uncertainty and fear of senior leaders who don't feel well equipped to go deep enough into their AI investments. And, of course, this pattern was accelerated and exacerbated by the arrival of tools like ChatGPT, Claude Code, and Copilot, among others, which serve as inspiration and reference points for companies and leaders looking to create their own applications and AI success stories.

The good thing is that I have seen so many well-structured, successful initiatives (before and during the generative AI era) that it seems clear to me that no investment can rely, at least not solely, on bottom-up, organic innovation. Don't get me wrong, there are many engineering-led organizations with amazing folks creating and implementing great ideas and figuring out the details along the way. But these are the exceptions, and they usually require a very high level of AI talent and an underlying company-level technology maturity that, unfortunately, we don't see in every organization.

On the other hand, I have observed (and, of course, been part of) quite a few failed initiatives: "solo" efforts from visionary leaders that become a waste of money despite their best efforts, AI projects with unrealistic stakeholder expectations, and unclear strategies that generate disconnected groups within the same company as a sort of competitive, rather than collaborative, exercise.

My favorite example for a few years was companies with very little technology and data maturity hiring one data scientist and hoping that this almost mythological creature would work their magic and tell them what to do by reading directly from the often-imperfect data sitting in more or less obsolete systems. When that idea failed, they would hire more folks and hope for a miracle. And nowadays, the increasing level of AI awareness is bringing great news, but also a fair dose of noise and brother-in-law-style opinions that don't contribute to the success of the company's efforts.

But coming back to the individual aspect of this art called *AI leadership*, two main streams summarize the concerns and hesitations I have heard from all these executives: how to work at scale to enable collaboration between relevant stakeholders, and how to structure an end-to-end approach that goes well beyond just processes and basic AI strategies and proactively brings tangible exercises and controls that help the rest of the organization (including technical departments) approach their share of the work in a well-synchronized and structured manner. And, in the meantime, there is the impostor syndrome of valuable leaders feeling overwhelmed and almost forced to innovate and prove something to someone, including higher-level executives, industry observers, and, of course, investors.

At a personal level, I feel there is so much untapped potential that companies and their leaders deserve a new way to approach their AI initiatives. Generative AI may include some degree of buzz and noise, but it also makes AI incredibly accessible to most companies out there. Honestly, I have never seen this level of adoption or even this degree of technical understanding of what technology can bring to those companies and their products and services. It is almost magical to join executive and board discussions and realize that people are starting to understand technology in a realistic and sustainable manner. The era of happy AI ideas seems to be evolving into something more mature, and all the stars are starting to align to create more value and more success stories.

But still, most companies struggle to materialize even the best ideas and use cases into something tangible. Regardless of the source, studies from organizations such as MIT or Gartner keep quantifying failure as something close to 80% of initiatives and projects worldwide. This is not new; I have references to articles talking about the same thing from a long time ago. The difference is that nowadays, more companies and leaders are taking the leap of faith and suffering the consequences. That said, in my opinion, these articles and studies often miss the point, as the probability of failure doesn't equal the probability of failure given that the AI project is properly structured, and that's the key.

$$P(A) \neq P(A/B)$$

... therefore, let's focus on "B" and skip too optimistic  
or too pessimistic analyses  
and expectations

*Figure 0.2 – A good example of why everyone should understand probabilistic and Bayesian methods*

But hey, even when some companies succeed, they often miss the point when it comes to dealing with experimentation and uncertainty, collaborating with third parties (e.g., technology providers), differentiating between noise and reality, connecting with the external context to incorporate trust and compliance aspects by design, or even expanding technical activities and connecting relevant data, AI, IT, and security stakeholders. If they had the proper techniques and processes to improve all of this, would we still be talking about 80%? Probably not.

For these and other reasons, I decided to park my hyper-technical books and courses for a bit and collaborate with Packt to bring something different (and, in my obviously biased opinion, quite special): an executive playbook that is accessible but reaches a proper level of technical detail; an easy-to-follow read that also includes tangible assets for leaders to use in their day-to-day activities; and, even more importantly, a book that doesn't end in high-level yada-yada and disappointment and that makes their time investment worthwhile, because no company

executive or leader has time to waste, and their professional success and evolution will heavily rely on the results of these big investments and AI initiatives.

If, after this long series of questions, unbacked opinions, and old-person stories, you are still willing to proceed, be reassured: this will help you. We will escape from theoretical frameworks and ideas, and I will bring every single technique and template that my brain (supported by cloud storage) can remember, directly extracted from my strategic, tactical, and technical experiences. And for that, we will use AI, and especially generative AI in the modern context, as the excuse for you to evolve your technology innovation toolkit and become that full-stack leader who can save the day, or at least make it better. Let's get started!

*Adrián González Sánchez*

*Your friendly neighborhood author*

## Who this book is for

The term "playbook" is a clear statement of intention. The goal is to provide readers with an easy-to-read yet meaningful and comprehensive overview of what AI, and generative AI in particular, means from a leadership perspective. This is an unapologetically executive-centric piece that combines direct explanations and practical recommendations with a well-structured thought process. It is designed to be immediately useful, offering concrete examples, templates, and reusable checklists that leaders can apply in their day-to-day decision-making.

From that perspective, this book is a strong fit if you belong to one of the following groups:

- **C-suite folks** and board members seeking strategic clarity on AI.
- **Senior leaders and decision-makers** responsible for transformation, technology, or business performance.
- **Aspiring executives** preparing to lead in an AI-enabled organization.
- **Other professionals and leaders** who want to strengthen their AI leadership skills and strategic judgment.

So yes, this playbook is not about becoming a data scientist or an AI engineer. It is about understanding how to lead, decide, and compete in a world where AI is becoming a core business capability. It is about influencing the way your AI initiatives are structured within your company.

Last but not least, you have chosen this book because you want real advice. You are not looking for basic explanations, but rather for an applied approach to understanding the value of AI in your company, without obvious statements or the classic introductory knowledge associated with "AI for executives" books. And you are doing this because you really want to learn and put in the work, so part of your role here will be to review the content, reflect on it, adapt the proposed templates, add your thoughts and notes, and revisit the chapters whenever you need to.

## What this book covers

*Chapter 1, Embracing Top-Down AI Approaches: The Leadership Imperative*, addresses the notion of modern leadership, your own executive engagement, how you spread your vision, and how AI is managed at scale. Your company's operating system must evolve, and this first chapter will help you identify the main decision points and the multiple levels of management you need to lead from the top to make sure your AI strategy becomes an operationalized, successful reality.

*Chapter 2, From Turing to ChatGPT: A Brief History and Context of AI*, explores previous cycles of hype and disappointment, but it also examines how time can indicate a technology's maturity and how that maturity affects investment value, risk, and uncertainty. Understanding these factors will help you look beyond the excitement surrounding AI and focus on the technologies most relevant to your highest-priority projects. As with many things in life, a historical perspective helps you form an informed opinion, which is an essential capability for any company leader.

*Chapter 3, Exploring Key AI Terminology and Concepts*, is very important. Most of the people out there will tend to skip any kind of pre-ChatGPT knowledge. And then they will mix terms like machine learning and generative AI. And what's most important, they will feel lost and use their mental energy to hide the lack of knowledge and pretend they are understanding what their technical teams are talking about. And that's the wrong way to approach AI leadership. You don't need to be an expert, but you certainly need to have a baseline knowledge that will allow you to ask the good questions when decisions such as project sponsorship, resource allocation, and use case prioritization must be made. This chapter equips you with all the non-Generative AI knowledge you need to understand the rest of AI technologies.

*Chapter 4, Understanding the New Era: Generative AI and Agentic AI*, explores the technical fundamentals behind generative AI and the agentic era. Explained from the perspective of a topic that is continuously evolving but that leverages a series of foundations that remain the same, and from the point of view of a company executive or leader, this chapter will conclude the most didactic section of this book (i.e., Chapter 2 through Chapter 4, spanning the end of *Part 1* and the start of *Part 2*) and solidify your technical understanding of the matter, before coming back to exec-led activities in the context of artificial intelligence.

*Chapter 5, Crafting an AI Strategy: From Vision to Execution*, allows you not only to define the main areas of your AI strategy, but also to evaluate the current level of maturity of your organization, and to run a maturity and gap assessment that will help you track the improvements of your AI strategy over time. This was one of the most challenging chapters for me as an author, because I wanted to bring something that goes beyond typical AI strategy discussions, but also a very rewarding one, as it will set the pillars of your next AI strategy.

*Chapter 6, AI-First Approaches: New Business Models and Products*, is based on the idea of frontier innovation. If you frame your AI strategy properly and your organization is ready to implement it, AI will give you an advantage that enables new kinds of business models and opportunities. This goes beyond the classic "saving money, generating revenue" mantra from management consulting firms. This type of innovation is really about transforming the nature of your business, at both external (e.g., customer-facing products) and internal (AI-enabled productivity, new strategic priorities, creating new ventures) levels. And of course, as with any other cycle of transformation, it will require some degree of organizational change and the corresponding change management programs. This chapter combines all of this.

*Chapter 7, Building and Scaling AI Teams: Roles, Reskilling, and Hiring*, explores all talent-related things in the context of AI. Run away from the classic "AI talent is scarce and we cannot find it" sentence. That is usually a consequence of more structural issues. This chapter deals with hiring new employees and training your existing workforce. You will also understand the evolution of the main roles, especially after the arrival of generative AI tools and technologies. And you will explore how to address the impact of AI on the future of work to maintain a healthy dynamic with your teams and employees. This seventh chapter can have a huge impact on how your AI strategy becomes a reality, as people and their skills are the most important factor.

*Chapter 8, Navigating the AI Ecosystem: Tools and Vendor Evaluation*, illustrates how the AI ecosystem is right now a mix of vendors, tools, technologies, sales folks, and work-in-progress capabilities that makes it very difficult for anyone to differentiate noise from reality. This chapter will help you differentiate the main categories of tools that will impact your AI projects (from internal productivity to dedicated cloud and AI platforms) and to properly assess vendors while leveraging the most suitable set of tools and technologies available.

*Chapter 9, Responsible AI Implementation: Ethics, Governance, and Accountability*, shows how Responsible AI approaches are nowadays a technical, societal, and reputational must for any AI adopter organization. Except if you are planning to lead your AI initiative without ethical limits (I highly recommend you not to), this chapter will help you first understand what an ethical implementation is, your role as the AI leader, and how to make it happen while enabling continuous innovation.

## **To get the most out of this book**

Readers should be leaders, managers, decision-makers, or professionals preparing to take on leadership responsibilities in an AI-enabled organization. No technical AI background is required, but curiosity about artificial intelligence and a willingness to engage with its business and technical dimensions will be valuable. Readers will get the most from this book by actively applying its frameworks, checklists, and exercises to their own organization, teams, and AI initiatives.



## Disclaimer on images

Some images in this title are presented for contextual purposes, and the readability of the graphic is not crucial to the discussion. You may refer to our free graphic bundle to download the images.

## Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here: <https://packt.link/gbp/9781806388493>.

## Conventions used

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

**Bold:** Indicates a new term, an important word, or words that you see on the screen. For instance, words in menus or dialog boxes appear in the text like this. For example: "**Financial operations (FinOps)** teams should see, understand, and discuss generative AI costs in near-real time."

### Note

**Warnings** or important notes appear like this.

### Tip

Tips and tricks appear like this.

## Get in touch

Feedback from our readers is always welcome.

**General feedback:** If you have questions about any aspect of this book or have any general feedback, please email us at [customer@packt.com](mailto:customer@packt.com) and mention the book's title in the subject of your message.

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

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

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

## Free benefits with your book

This book comes with free benefits to support your learning. Activate them now for instant access (see the "*How to Unlock*" section for instructions).

Here's a quick overview of what you can instantly unlock with your purchase:



### DRM-Free PDF Version

Download DRM-free PDF and ePub copies of this book.



### 7-Day Packt Library Access

Get 7-day unlimited access to 8,000+ books and videos. No credit card required.

Available for first-time Packt+ trial users only.



### Next-Gen Reader Access

Read this book on Packt Reader with progress sync, dark mode and note-taking.

## How to Unlock

Scan the QR code (or go to [packtpub.com/unlock](https://packtpub.com/unlock)). Search for this book by name, confirm the edition, and then follow the steps on the page.



UNLOCK NOW

*Note: Keep your invoice handy. Purchases made directly from Packt don't require one*

## Share your thoughts

Once you've read *Generative AI for Leaders Playbook*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback.



<https://packt.link/r/1806388499>

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

# Part 1

## Building Your AI Leadership Foundation

Before leaders can make good AI decisions, they need the right leadership model, historical perspective, and technical vocabulary. *Part 1* establishes that foundation: how to lead AI adoption from the top, what previous waves of AI can teach us about hype, maturity, and risk, and the essential concepts leaders need to communicate effectively with technical teams and make informed decisions about investments, resources, and use cases.

This part contains the following chapters:

- *Chapter 1, Embracing Top-Down AI Approaches: The Leadership Imperative*
- *Chapter 2, From Turing to ChatGPT: A Brief History and Context of AI*
- *Chapter 3, Exploring Key AI Terminology and Concepts*



# 1

## Embracing Top-Down AI Approaches: The Leadership Imperative

Let me tell you something upfront: your role as a company's executive, leading the clearly defined exec-led task to enable your organization's AI adoption, is not optional anymore. While most of the organizations out there are still trying to figure out the optimal way to implement and leverage AI, and more concretely, generative AI, now is the right time for you to embrace a *new leadership approach* in which you can set the pace and scope of your AI innovations, directly from the top.

But for that to happen, you need to understand and embrace the key fundamentals, including what AI represents and its impact on your business, the different levels of management (very important if you want to translate your AI strategy into actionable streams at multiple levels), the main operational methodologies, the notion of an end-to-end AI lifecycle, and the realistic applications of AI depending on your company's context (e.g., access to talent, tools, and overall level of tech maturity).

While all these concepts sound a bit far from the actual day-to-day of a company executive or decision maker, that's the precise goal of this chapter. We will go beyond just high-level "101" information and the general AI hype and help you upskill to a certain level of knowledge that converges with the expertise from your technical teams. This is important because, without that shared level of understanding, you cannot credibly challenge assumptions, prioritize investments, or separate real opportunities from hype. In other words, you would be delegating strategy, not just execution. You are now the **leader behind your company's AI initiative**, and you shouldn't delegate that mix of power and responsibility!

The problem, in short, is that most organizations are experimenting with AI, but very few have the executive-level direction required to turn that experimentation into durable business value. The solution starts with you, and this chapter will give you the initial toolkit. By the end of this chapter, you will be able to explain what AI and generative AI really mean for your business, create your first AI policy, design a multidisciplinary AI committee, align AI use cases with your corporate strategy, and evaluate the strategic, tactical, and technical levels of AI management in your organization.

Specifically, we will cover the following topics:

- Understanding what artificial intelligence means today
- Assessing generative AI for innovation and executive engagement
- Understanding the levels of AI management

#### Note

##### Download the PDF version of this book

Your purchase includes a DRM-free PDF copy of this book, the code bundle, and a range of exclusive benefits. To unlock everything, follow the *Free benefits with your book* section in the *Preface*.

## Understanding what artificial intelligence means today

It sounds like a basic and obvious question, but the answer is less obvious and more complex than one may think. Today, when everyone is talking about AI, are we sure that our understanding of what artificial intelligence means and includes is still correct? Do you have a specific definition in mind? Is there a single and universal one? Depending on the sources, you will have multiple definitions of what AI means, and even a variable scope of technologies that are considered artificial intelligence. Let's start with the technical fundamentals.

The core aspect of AI is described by industry experts like Alan Turing, John McCarthy, and many others as the ability to *enable machines to perform tasks that typically require human cognitive capabilities*. Note that Turing's famous test, for example, was about observable behavior, whether a machine could convincingly imitate a human in conversation, rather than about literally replicating human intelligence, a nuance that still matters when you assess what today's AI can and cannot do. Obviously, the level of capabilities has increased over time, with more advanced and powerful models, but the original concept remains the same.

There is also an interesting definition which connects AI with the notion of **digital intelligence**. IVADO, a Canadian AI research and knowledge transfer consortium based in Montreal, defines digital intelligence as *the set of tools and methodologies that, in combining the collection and harnessing of data with the design and use of models and algorithms, facilitates, enriches and supports decision-making* (<https://ivado.ca/en/digital-intelligence-to-move-from-data-to-decision/>). This refers to all data-driven technologies that may use AI or others to accomplish decision-making tasks. *Figure 1.1* illustrates how these elements, i.e., data, models and algorithms, and decision-making, come together in IVADO's definition of digital intelligence:




Figure 1.1: IVADO's definition of digital intelligence

We will go into the details of these technologies, i.e., data, models and algorithms, and decision-making, in the upcoming chapters, but the definition includes everything from business intelligence with dashboards, to operational research for optimization topics, and the entire family of machine/deep/reinforcement learning that fuels modern AI applications.

Even if one of your motivations behind picking this book up is learning about all AI types, you are primarily here because of the **generative AI boom**. You want to make sense of this new wave of **large language models (LLMs)** and AI agents. In this book, we will explore their impact on your activities, starting with your role and involvement as an executive and leader in your organization.



## Assessing generative AI for innovation and executive engagement

Even if generative AI techniques were already an important topic in the AI research industry before the release of ChatGPT by OpenAI in November 2022, it was the exact moment that brought generative AI capabilities to mainstream consumer awareness and made this technology relevant for most of the companies out there, including yours. To be clear, AI was already well established in the enterprise world; think of recommendation engines, fraud detection systems, or voice assistants that had been running in production for years, but ChatGPT put generative AI in everyone's hands. It makes sense: powerful technologies becoming accessible via regular language, enabling millions of people to try and experiment every day in a way that wasn't possible with other traditional, developer-oriented AI technologies.

Currently, the arrival of generative AI and the increasing level of general awareness and technology adoption are making the technical exploration simpler, but the reality is that the overall level of actual knowledge and understanding of the AI fundamentals, among both leadership teams and the broader workforce, is still far from ideal. That factor impacts not only your technical activities in the company, but also how the rest of the organization perceives AI, its value, and the level of realism of their expectations and proposed projects and use cases.

Even more important. Generative AI is now everywhere, and that explosive presence brings new requirements for you as a leader or C-suite member to not only upskill but also bring a renovated vision in both internal and external circles. Think about it. From Davos to the geo-political scene, from industry events to any internal executive meeting: generative AI is omnipresent and almost a mandatory requirement from investors and the public perspective.

This change brings new requirements for you, and the **inevitable impostor syndrome**. If you feel that everything goes so fast that it is impossible to catch up, and that you or even the rest of the leadership team are not equipped to act and lead the AI adoption, you are not alone. Also, if you perceive that everyone is talking about AI and you suspect that there are actual impostors in the room, you may be right. The anxiety to become relevant, have an opinion, showcase a certain level of knowledge, and the need to speak up during any AI-related meeting is provoking scenarios in which it is difficult to differentiate knowledge from noise, and valuable folks from just talkers.

The good news is that you don't need to become a machine learning engineer to overcome that feeling. Focus your personal upskilling on the short list of concepts that actually drive executive decisions: what LLMs can and cannot do reliably, why they hallucinate, the difference between adapting a model via prompts, retrieval techniques, or fine-tuning, where the main costs come

from, and how AI systems are evaluated and monitored once they reach production. We will cover all of these in this book, at the right level of depth for your role.

As for separating knowledge from noise in a meeting, a few simple questions work remarkably well: How do we measure whether this works? What happens when the model is wrong, and who is accountable? What data does it need, and are we allowed to use it? What will it cost to run at scale? People with real expertise will answer with specifics, trade-offs, and honest unknowns, while the talkers will answer with more buzzwords. You don't need to have all the answers yourself. Asking the right questions is precisely your job as a leader.

That AI knowledge gap, along with the **inherent nature of AI projects** (e.g., uncertain, difficult to estimate, still hard to prove its ROI), makes AI adoption a risky investment, and some sort of leap of faith in which your company *hopes* to get the classic cost-saving, revenue generation, or differentiation advantages. However, would it be possible without the right strategies?

So perhaps becoming the **main champion for AI adoption** in your company is not enough. You will need to leverage your formal management power and your natural leadership skills to influence, motivate, and guide other AI champions in the organization, so the culture spreads progressively, and general topics such as upskilling, hiring, and collaboration get adapted to the AI era in a natural way.

For all this to happen, you will need to create a system that combines a general AI vision with multiple levels of AI-oriented management that then connect with the day-to-day operations that will bring your AI developments to life.

As a business leader, you will need to develop and share the **AI vision for your organization**. That means defining how your company will adopt AI, what the main goals of that adoption are, how the AI use cases will support the overall strategy of the firm, and how you intend to make this actionable, as a top-down exercise that reaches all the levels of the hierarchy.

As part of your executive generative AI playbook, here are the assets and templates that will help you consolidate your company's AI vision.

## Creating your first AI Policy

As you may be aware, an **AI Policy** is a company-wide formal internal framework that defines how artificial intelligence tools and systems may be used across the organization. It establishes consistent rules, ethical standards, and governance processes to ensure AI is used responsibly, legally, securely, and in alignment with the company's values and risk tolerance.

The policy can apply to all employees, contractors, and business units, providing shared expectations and accountability for AI-related activities. It is often made available as an internal website or simply a PDF document for acceptance and signature.

Besides a brief explanation of the general AI vision, some of the key areas that you may add to your AI policy are:

- **Purpose and scope:** This section defines why the AI policy exists and who and what it applies to. It clarifies the policy's objectives (such as enabling responsible innovation, managing risk, and ensuring regulatory compliance) and specifies the scope of AI systems, tools, and use cases covered. As a governance framework template for internal AI usage, it sets boundaries for acceptable use across departments while allowing room for controlled experimentation and business-specific adaptations.
- **Guiding principles:** This section outlines the core values that guide all AI use within the organization. It typically includes principles such as transparency, fairness, accountability, privacy, security, and human oversight. Provided as a draft list with blanks, it allows the company to tailor the principles to its industry, regulatory environment, and corporate values, ensuring ethical considerations are embedded in every AI initiative.
- **Approved and prohibited use cases and tools:** This section specifies what AI may and may not be used for. It includes pre-filled categories of approved use cases (e.g., productivity support, data analysis, customer service augmentation) and prohibited uses (e.g., unlawful surveillance, discriminatory decision-making), with space for customization. The goal is to reduce ambiguity, prevent misuse, and help employees quickly determine whether a proposed AI use is allowed or requires escalation.
- **Compliance and security requirements:** This section details the legal, regulatory, and security obligations associated with AI use. It references applicable laws, standards, and internal policies, and includes requirements for data protection, anonymization, access controls, and secure model usage. By setting minimum compliance and security expectations, it helps mitigate legal, reputational, and cybersecurity risks associated with AI systems.
- **Roles, responsibilities, and enforcement:** This section defines who is responsible for what in AI governance. It outlines key roles (e.g., policy owner, AI governance committee, IT/security, business users), their responsibilities, and decision-making authority. It also describes enforcement mechanisms, escalation paths, and review frequency to ensure the policy remains effective, up to date, and consistently applied across the organization.

If your company has never created one, releasing and sharing the AI policy with the rest of the organization will be a great way to share that vision and make sure everyone adheres to and signs the policy. If there is already one in place, the AI policy becomes a living document that can be updated periodically to include new examples, recommendations, and requirements.

Here is an example you can reuse and adapt. Feel free to add specific considerations related to your own AI vision and your company's context. Keep in mind that this template is a starting point, not legal advice: before bringing any version of it into effect, make sure it is reviewed by your corporate legal and compliance teams. And if you are already familiar with AI policies, feel free to skim the template and jump ahead to the next section.

**Tip****Company AI Usage Policy**

**Document Status:** Draft / Approved

**Effective Date:** [DD/MM/YYYY]

**Policy Owner:** [Role / Department]

**Review Frequency:** [e.g., Annual / Semi-Annual]

**Last Review Date:** [DD/MM/YYYY]

**1. Purpose and Scope****1.1 Purpose**

This AI Usage Policy establishes a company-wide framework for the responsible, ethical, and secure use of artificial intelligence (AI) technologies. The policy is intended to:

- Share the general AI vision of the organization
- Enable innovation and productivity through appropriate AI use
- Ensure compliance with applicable laws, regulations, and contractual obligations
- Protect company, customer, and employee data
- Mitigate legal, ethical, operational, and reputational risks
- Align AI use with the company's values and risk appetite

**1.2 Scope**

This policy applies to:

- All employees, contractors, consultants, and temporary staff
- All departments, subsidiaries, and business units
- All AI systems, tools, models, and services, including generative AI, machine learning, and third-party AI platforms

This policy covers AI used for internal operations, customer-facing activities, and decision-support processes.

### 1.3 Definitions

- Artificial Intelligence (AI): Systems or tools that perform tasks requiring human-like intelligence, such as learning, reasoning, content generation, or decision-making.
- AI Tool/System: Any software, model, or service that uses AI techniques.
- High-Risk AI Use: AI applications that may significantly impact individuals, customers, employees, or business-critical decisions.
- Personal Data: Any information relating to an identified or identifiable individual, as defined by applicable data protection laws.
- [Other Definition] - [Add any other relevant definition]

## 2. AI Vision and Strategic Direction

### 2.1 AI Vision

The company's AI vision defines how artificial intelligence will be leveraged to support long-term business strategy, innovation, and responsible growth. AI is viewed as a strategic enabler to augment human capabilities, improve decision-making, and deliver sustainable value while adhering to ethical, legal, and security standards.

### 2.2 Main Goals

The primary goals of the company's AI strategy include:

- Enhancing operational efficiency and productivity
- Improving quality, consistency, and speed of decision-making
- Enabling innovation in products, services, and customer experiences
- Strengthening risk management, compliance, and security capabilities
- Ensuring responsible, ethical, and transparent use of AI technologies

### 2.3 Key Stakeholders

The successful implementation of the AI vision involves collaboration among the following stakeholders:

- Board of Directors and Executive Leadership

- Business Unit Leaders and Product Owners
- Employees and End Users
- IT, Data, Security, Legal, and Compliance Functions
- Customers, Partners, and Vendors
- Regulators and External Auditors (where applicable)

#### **2.4 Expected Gains and Value Creation**

The company expects AI adoption to deliver measurable benefits, including:

- Cost savings and efficiency gains through automation and augmentation
- Improved customer satisfaction and service quality
- Better insights derived from data-driven analysis
- Reduced operational and compliance risks
- Increased competitiveness and long-term business resilience

#### **2.5 Mid- to Long-Term Plans**

The company's AI roadmap will be reviewed and updated periodically and may include:

##### **Mid-Term (1–3 years):**

- Establishing mature AI governance and risk management processes
- Scaling approved AI use cases across business units
- Investing in employee AI skills and training programs
- Standardizing AI tools, platforms, and data practices

##### **Long-Term (3–5+ years):**

- Embedding AI capabilities into core business processes and products
- Advancing responsible and explainable AI practices
- Leveraging AI for strategic differentiation and innovation
- Continuously aligning AI strategy with evolving regulations and technologies

### 3. Guiding Principles

All AI usage within the company must align with the following principles:

- **Transparency:** AI use should be clearly disclosed where appropriate, and outputs should be explainable to a reasonable degree.
- **Fairness & Non-Discrimination:** AI systems must be designed and used to avoid bias and discriminatory outcomes.
- **Accountability:** Human oversight must be maintained, with clear ownership of AI-driven outcomes.
- **Privacy and Data Protection:** AI use must respect data protection laws and internal privacy standards.
- **Security & Resilience:** AI systems must be used securely and protected against misuse or unauthorized access.
- **Human Oversight:** AI should augment, not replace, human judgment in critical decisions.
- **[Custom Principle]:** [Add organization-specific principle]

### 4. Approved and Prohibited AI Use Cases

#### 4.1 Approved AI Use Cases (Examples)

The following categories are generally permitted, subject to compliance with this policy:

- Productivity and efficiency support (e.g., drafting content, summarizing documents)
- Data analysis and insights generation using approved datasets
- Customer support augmentation (non-decisive, human-reviewed)
- Software development assistance (e.g., code suggestions), using approved AI tools only — see Section 7.2 for the restrictions that apply to public tools
- Training and learning support

**Note:** Approval may still be required for high-risk or large-scale deployments.

#### 4.2 Prohibited AI Use Cases

The following uses are strictly prohibited unless explicitly approved by the AI Governance Body:

- Processing confidential, sensitive, or personal data in unapproved AI tools
- Automated decision-making with legal or significant impact without human review
- Discriminatory profiling or surveillance of individuals or groups
- Use of AI to generate misleading, deceptive, or unlawful content
- Training AI models on proprietary or third-party data without authorization

#### **4.3 Use Case Approval Process**

- New or high-risk AI use cases must be submitted for review to: [Committee / Role]
- Risk assessments and documentation may be required prior to approval

### **5. Compliance and Security Requirements**

#### **5.1 Legal & Regulatory Compliance**

AI usage must comply with all applicable laws and standards, including but not limited to:

- AI, data protection, and privacy laws (e.g., EU AI Act, GDPR, CCPA, or local equivalents)
- Employment, consumer protection, and intellectual property laws
- Industry-specific regulations

#### **5.2 Data Handling and Anonymization**

- Personal or sensitive data must not be input into AI systems unless explicitly approved
- Data should be anonymized or pseudonymized where feasible
- Only minimum necessary data may be used

#### **5.3 Security Controls**

- Use only company-approved AI tools and platforms



- Follow access control, authentication, and logging requirements
- Report suspected AI-related security incidents immediately
- Assess and mitigate AI-specific attack vectors, such as prompt injection, before deploying LLM-powered tools that process untrusted content

## **6. Roles and Responsibilities**

### **6.1 Employees and Users**

- Use AI tools in accordance with this policy
- Complete required AI training
- Escalate concerns, risks, or incidents

### **6.2 Managers**

- Ensure team compliance with the policy
- Approve low-risk AI use cases where delegated

### **6.3 AI Governance Body / Committee**

- Review and approve high-risk AI use cases
- Maintain and update the AI policy
- Monitor regulatory developments

### **6.4 IT / Security / Legal**

- Provide approved AI tools and technical controls
- Conduct risk and compliance assessments
- Support investigations and incident response

## **7. Internal Use of AI Tools and Data Handling**

### **7.1 Approved vs. Unapproved AI Tools**

Employees may only use AI tools that are explicitly approved by the company for business purposes ("Approved AI Tools"). Approved tools may include internally hosted AI systems or vetted third-party solutions under contractual and security controls.

The use of unapproved, public, or free AI tools (e.g., publicly available generative AI platforms) for company-related work is restricted and subject to the requirements below.

## **7.2 Use of Personal, Confidential, and Sensitive Information**

### **Approved / Private AI Tools:**

- Personal data, confidential information, and internal business data may only be processed if the tool is approved for such use and appropriate safeguards are in place.
- Data processing must comply with applicable data protection laws and internal privacy policies.
- Only the minimum necessary data should be used for the intended purpose.

### **Public or Free AI Tools:**

- Employees must not input personal data, sensitive data, confidential information, trade secrets, source code, or proprietary business information.
- Information entered into public AI tools should be treated as potentially stored, reused, or disclosed by the provider.
- Only generic, non-identifiable, non-confidential information may be used.

## **7.3 Prohibited Data Inputs (All Tools)**

The following data must never be entered into any AI tool unless explicitly approved:

- Personal data of customers, employees, or third parties
- Special category or sensitive personal data (e.g., health, biometric, financial)
- Confidential business information or intellectual property
- Credentials, passwords, or security keys

## **7.4 Output Review and Human Oversight**

- AI-generated outputs must be reviewed by a human before use or distribution.
- Users are responsible for validating the accuracy, legality, and appropriateness of AI outputs.

- AI outputs may include hallucinations, such as confident-sounding but incorrect or fabricated content, so facts, figures, citations, and code must be verified before being relied upon
- AI-generated content must not be presented as human-generated where disclosure is required.

### **7.5 Recordkeeping and Transparency**

- Material AI-assisted work may need to be documented or disclosed internally.
- High-risk or customer-facing AI use must follow additional transparency and documentation requirements.

### **8. Training and Awareness**

- Mandatory AI policy and responsible-use training may be required
- Additional training may be required for high-risk or specialized AI use
- [Add examples of internal or external AI trainings in place]

### **9. Enforcement and Violations**

Failure to comply with this policy may result in:

- Suspension of AI access
- Disciplinary action up to and including termination
- Legal or contractual consequences
- Violations or suspected misuse must be reported to: [Contact / Function]

### **10. Monitoring, Review, and Updates**

- This policy will be reviewed at least [annually]
- Updates may be made to reflect legal, technological, or business changes
- Employees will be notified of material changes

### **11. Exceptions**

Any exceptions to this policy must be:

- Documented

- Risk-assessed
- Approved by: [Authority]

### Acknowledgement

I acknowledge that I have read and understood the AI Usage Policy and agree to comply with its requirements.

Name: \_\_\_\_\_

Signature: \_\_\_\_\_

Date: \_\_\_\_\_

## Defining your new AI committee

One of the most important decisions and regular tasks of AI leaders is to **define a proper structure and human group for the AI committee**. As you may know, an AI Committee is a multidisciplinary governance body responsible for overseeing the strategic, ethical, and risk-aware use of artificial intelligence across the organization. It brings together executive leadership and key functional representatives such as business, technology, legal, risk, security, and HR, to ensure AI initiatives align with company strategy, regulatory obligations, and risk appetite. The committee provides direction, approves or escalates higher-risk AI use cases, sets governance standards, and promotes responsible innovation through a structured, risk-based approach.

As part of your playbook, here is your own **AI committee checklist** so you can design your overall structure and keep in mind all key steps:

### Tip

#### AI Committee Checklist

1. **Define the Committee's Mandate Before Selecting Members:** Before appointing individuals, clearly define the AI committee's purpose and decision rights. The committee should be positioned as a strategic and governance body, not only a technical forum. It should have the following responsibilities:

#### (Committee Activities Checklist)

- ☐ Setting AI strategy and priorities
- ☐ Approving or overseeing high-risk AI use cases
- ☐ Defining and maintaining AI governance standards
- ☐ Monitoring regulatory, ethical, and operational risks

☐ Acting as an escalation and decision-making forum

☐ [other responsibilities and goals]

**Add your notes here:**

---



---



---



---

Remember, a clear mandate helps ensure the right skills are selected and prevents overlap with existing committees. It also sets the scope and approach of the committee activities, allowing you to choose or nominate the best potential members.

**2. Build the Committee Around Complementary Skill Sets, Not Titles:** The effectiveness of an AI committee depends on the combination of skills, not seniority alone. The goal is to balance innovation, risk, and business value.

**(Areas and Skills Mix Checklist)**

- ☐ **Executive leadership:** Strategic alignment, investment decisions, and authority
- ☐ **Business and product expertise:** Understanding real use cases and value creation
- ☐ **Technology and data:** AI capabilities, limitations, and system architecture.

Be specific here: you will want someone who can evaluate models and their benchmarks (e.g., a senior AI/ML engineer or applied scientist), someone who understands data pipelines and MLOps in production, and someone with enough hands-on experience to challenge vendor claims and not just a senior title from the IT organization

- ☐ **Legal and compliance:** Regulatory interpretation, liability, and contractual risk
- ☐ **Security and privacy:** Data protection, cybersecurity, and resilience
- ☐ **Risk management or internal audit:** Structured risk assessment and controls
- ☐ **HR or people functions:** Workforce impact, training, and change management
- ☐ [other company areas and skills]

**Add your notes here:**

---



---



---

Remember to avoid overloading the committee with specialists from one domain. The diversity of perspectives will be essential for balanced decision-making. These multidisciplinary groups can include people with technical, executive, ethical, and other backgrounds. Also, try to leverage people's genuine interest and excitement, and avoid falling into the impostor trap (i.e., bringing to the committee self-proclaimed experts with strong personal goals instead of willingness to contribute and collaborate).

**3. Identify Executive Sponsors Across Key Company Areas:** To ensure authority and adoption, the AI committee should have visible executive sponsorship from all relevant parts of the organization, such as:

**(Executive Sponsor Checklist)**

- ☐ Technology / IT: [add examples of names and roles]
- ☐ Legal and Compliance: [add examples of names and roles]
- ☐ Risk or Internal Audit: [add examples of names and roles]
- ☐ Core Business Units: [add examples of names and roles]
- ☐ HR and Data Governance: [add examples of names and roles]
- ☐ [add other company areas, when relevant]

**Add your notes here:**

---



---



---



---

Each of these sponsors should commit to at least:

- Act as a champion for AI adoption and governance within their function
- Ensure alignment between committee decisions and functional priorities
- Provide resources and support for implementation
- Remove any potential roadblocks and internal politics
- Identify new skills and training needs within their area

This shared sponsorship model prevents AI governance from being seen as "owned" by only one department and increases organizational buy-in, which is one of the main risks of leading the AI initiative from a specific area of the company or level of management.

**4. Appoint Operational Champions to Bridge Strategy and Execution:** In addition to executives, include key AI champions who are closer to day-to-day operations. You can define the AI champions' function as something complementary to the employees' official role, or hire specifically for this role.

**(AI Champions Checklist)**

- ☐ Understand how AI is actually being used or proposed
- ☐ Translate policy into practical guidance
- ☐ Act as two-way communication channels between teams and leadership
- ☐ Identify, discuss, and escalate potential risks and concerns
- ☐ Coach and potentially train other employees
- ☐ [add other responsibilities, depending on the scope]

**Add your notes here:**

---



---



---



---

Champions help ensure the committee's decisions are implementable and grounded in operational reality. You can also find examples from organizations defining them as Responsible AI Champions: Microsoft, for instance, has publicly described its internal network of *Responsible AI Champs* as part of its responsible AI program (see <https://blogs.microsoft.com/on-the-issues/2021/01/19/microsoft-responsible-ai-program/>) and similar communities exist in other large technology companies.

**5. Embed a Risk-Based Approach as a Core Operating Principle:** The committee should explicitly adopt a risk-based approach to AI governance, recognizing that not all AI use cases carry the same level of risk.

**(Risk-based Approach Checklist)**

- ☐ Classifying AI use cases by risk level (e.g., low, medium, high) based on internal policies or regulatory requirements

- ☐ Focusing committee review and approval on higher-risk use cases, then replicating and automating for others
- ☐ Defining clear escalation criteria for sensitive or regulated use cases
- ☐ Aligning risk assessment with business impact, data sensitivity, and regulatory exposure
- ☐ Validating new use cases and initiatives against the overall AI policy
- ☐ Maintaining an updated inventory of applicable AI, data protection, and security regulations and standards
- ☐ [add other items, depending on the adjusted approach]

**Add your notes here:**

---

---

---

---

This approach enables innovation at speed for low-risk use cases while ensuring appropriate scrutiny where it matters most. Make sure your employees understand the value of this approach, so they don't perceive the additional tasks as a burden and stopper for their day-to-day activities.

**6. Align Committee Processes With Existing Governance Structures:** The AI committee should integrate with, not duplicate, existing governance bodies such as risk committees, data governance forums, or security councils. This is key due to the convergence of AI as a technology with all these areas, but also a great lever for you to accelerate the operationalization of the AI committee.

**(Governance Alignment Checklist)**

- ☐ Define clear interfaces between committees and escalation paths
- ☐ Reuse existing risk, compliance, and approval processes where possible
- ☐ Ensure reporting lines to executive leadership or the board
- ☐ Alignment reduces friction and strengthens accountability

**Add your notes here:**

---

---



---

---

In this case, your role as an executive AI leader will contribute to helping align the different parties to make sure the AI committee brings complementary value to the current structure.

**7. Set Clear Review Cadence and Decision Transparency:** Establish a regular meeting cadence and documented decision-making processes. As with any other type of governance, recurrence and habit makes the process more solid and help get people's attention.

**(Committee Process Checklist)**

- ☐ Review AI strategy and risk posture periodically
- ☐ Track approved AI use cases and key risks
- ☐ Document decisions, rationales, and conditions for approval
- ☐ Publish and share illustrative past cases
- ☐ Collect feedback from the different areas (e.g., enabling a single email address or an anonymous feedback survey)

**Add your notes here:**

---

---

---

---

This level of structure and transparency will help you build trust across the organization, and it supports auditability and regulatory readiness for the future.

**8. Establish and Maintain a Centralized Inventory of Tools and Technologies:** The AI Committee should be responsible for overseeing the systematic collection, cataloging, and maintenance of an inventory of tools and technologies related to AI, data, security, and IT across the organization. This inventory supports transparency, risk management, and informed decision-making, and serves as a foundation for effective AI governance.

**(Tools Catalog Checklist)**

- ☐ Maintain a centralized register of AI, data, analytics, automation, and generative AI tools in use or under evaluation

- ☐ Distinguish between approved, restricted, and prohibited tools
- ☐ Document key attributes for each tool, including purpose, business owner, data types processed, deployment model (internal, cloud, third party), and risk level
- ☐ Identify tools that interact with personal, sensitive, or confidential data
- ☐ Track security, privacy, and compliance assessments associated with each tool
- ☐ Align the inventory with vendor management, IT asset management, and data governance processes
- ☐ Regularly review and update the catalog to reflect new tools, decommissioned systems, and changes in risk profile

**Add your notes here:**

---

---

---

---

A well-maintained inventory enables the committee to apply a consistent, risk-based approach, reduce shadow AI usage, and respond efficiently to audits, incidents, or regulatory inquiries.

#### **9. Leverage External Experts and Consultants to Complement Internal Expertise:**

The AI Committee should proactively engage external experts and consultants when specialized knowledge or independent perspectives are needed to support informed decision-making. External input helps address skill gaps, challenge internal assumptions, and ensure alignment with evolving best practices and regulatory expectations.

##### **(External Involvement Checklist)**

- ☐ Engage external experts for emerging or complex topics such as advanced AI models, explainability, bias mitigation, or sector-specific regulation
- ☐ Use independent advisors to provide objective risk, ethics, or compliance assessments for high-impact AI initiatives
- ☐ Involve legal, regulatory, or standards experts to interpret new or evolving AI-related laws and guidance
- ☐ Invite external specialists to support maturity assessments, audits, or benchmarking exercises

- ☐ Ensure knowledge transfer from consultants to internal teams to build sustainable internal capabilities
- ☐ Define clear scope, decision authority, and confidentiality requirements for external advisors

**Add your notes here:**

---

---

---

---

Remember to periodically reassess the need for external support as internal AI maturity increases, as you will certainly spot potential gaps and new needs. Incorporating this external expertise enables the AI Committee to broaden its perspective, strengthen its risk-based approach, and make more robust, future-proof decisions.

**10. Evolve the Committee as AI Maturity Grows:** Recognize that the AI committee is not static. As AI adoption and regulation evolve, you will need to periodically reassess what you need to continue making the AI committee as functional and useful as possible.

**(Committee Evolution Checklist)**

- ☐ Skill gaps and representation
- ☐ Decision scope and authority
- ☐ Risk tolerance and governance depth
- ☐ Potential periods (e.g., X months) for iterations
- ☐ Backlog with learnings and potential improvements

**Add your notes here:**

---

---

---

---

An adaptable committee structure ensures long-term effectiveness and relevance. Don't get stuck; lead any potential improvement based on data points and consensual agreements. That will also increase the sense of ownership and accountability of different committee members.

## Confirming your North Star for AI use cases

Before starting to identify potential AI use cases and applications with your teams, as an AI leader you need to confirm your company's "AI North Star". This is the strategic direction that will guide all AI technology investments. AI should never be adopted in isolation or driven solely by novelty. It must directly reinforce the company's overall business strategy and long-term competitive positioning. I have seen so many AI initiatives that are obviously misaligned with strategy, and they often result in fragmented pilots, limited ROI, and organizational fatigue. Consider two illustrative (and unfortunately common) patterns: a retailer whose data team builds an impressive demand-forecasting prototype that the merchandising area never adopts, because it was not connected to any commercial priority, or a bank that rushes out a customer-facing chatbot while its actual strategy is built on premium, high-touch human advice. Both are technically fine projects, and both are strategically pointless.

So how do you make sure your AI goals are aligned with the corporate strategy? To ensure alignment, AI use cases should be evaluated against the company's core objectives, such as how the company competes, where it creates value, what differentiates it from competitors, and which capabilities are mission-critical.

## Using Porter's Generic Strategies Framework

Michael Porter's **Generic Strategies Framework** provides a practical lens to test whether AI use cases truly support the company's competitive intent. This framework is generally used to define or analyze the key strategic choice of a company, but you will also leverage it for your AI purposes. *Figure 1.2* shows the classic representation of this framework, with the generic strategies positioned according to competitive scope and competitive advantage:




Figure 1.2: Porter's Generic Strategies

According to Porter, organizations typically compete through one (or a combination) of the following strategies:

- **Cost leadership:** This means competing by being the lowest-cost producer in the market.
- **Differentiation:** This means competing by offering unique products, services, or experiences.
- **Focus:** This means competing by serving a specific niche or segment better than anyone else.

At the same time, your AI initiative and the specific projects you bring to life can or cannot support each of those strategies. For example:

- **AI for cost leadership:** This includes AI use cases that focus on automation, efficiency, waste reduction, and operational optimization. Thinking about demand forecasting, process automation, predictive maintenance, cost-to-serve optimization, etc.
- **AI for differentiation:** This includes AI projects that intend to enhance customer experience, personalization, product innovation, or decision quality. For example:

recommendation engines, AI-driven design, intelligent customer support, advanced analytics.

- **AI for focus:** This could be about AI use cases that deepen understanding of the target segment and tailor offerings precisely via niche customer insights, specialized pricing models, segment-specific automation.

Before you apply this mapping, keep in mind two important nuances. First, AI for cost leadership carries a built-in paradox: running LLMs at scale can be expensive, so pursuing cost leadership with costly tools only works when the unit economics are clear. That means targeted automation with measurable savings, smaller or optimized models where they are good enough, and a disciplined eye on consumption costs. A well-known illustration is Klarna, which publicly reported that its generative AI assistant was handling a large share of its customer service conversations, performing work it had previously distributed across hundreds of agents. Klarna later rebalanced the model and brought human agents back for complex cases, which is itself a lesson about unit economics versus service quality.

Second, if everyone builds on the same foundation models from the same few providers, the models themselves stop being a differentiator, and any advantage based purely on "using AI" erodes fast. Lasting differentiation comes from your proprietary data, your workflows, and how deeply AI is embedded into the product experience. Think of how companies like Netflix or Duolingo differentiate through AI-powered experiences built on their own data and user base, rather than through exclusive access to models. The same logic applies to focus strategies: specialized players serving a single vertical, such as legal or healthcare, can outperform generalist tools precisely because their data and workflows are narrow and deep.

AI use cases that do not clearly reinforce at least one of these strategic positions risk diluting focus and weakening competitive advantage. But also, if your selected use cases focus on the wrong value proposition and they are not aligned with the overall strategy of the company, who is going to sponsor them? Would you support something that doesn't generate relevant impact and is not even related to your KPIs and end-of-year bonus metrics? I don't think so.

Use the following checklist to confirm that proposed AI initiatives align with your North Star:

**Tip****AI North Star Checklist**☐ **Strategic Alignment**

- Does this AI use case clearly support our chosen competitive strategy (cost leadership, differentiation, or focus)?

- Can leadership articulate why this use case matters to the business, not just what it does?
- Would abandoning this use case weaken our strategic position?

**Add your notes here:**

---

---

---

---

☐ **Value Creation**

- Is the expected value measurable in terms leadership cares about (revenue, margin, cost, risk, customer outcomes)?
- Does this use case improve a core value driver rather than a peripheral activity?

**Add your notes here:**

---

---

---

---

☐ **Capability Fit**

- Does this use case strengthen a capability we want to be best-in-class at?
- Is it aligned with our operating model, talent, and culture?

**Add your notes here:**

---

---

---

---

☐ **Prioritization and Focus**

- Does this use case compete for resources with higher-impact initiatives?
- Are we saying "no" to other AI ideas to protect strategic focus?

**Add your notes here:**

---

---

---

---

☐ **Sustainability**

- Will this use case create lasting advantage, or is it easily replicable by competitors?
- Does it build data, models, or learning effects that compound over time?

**Add your notes here:**

---

---

---

---

In subsequent chapters, you will use other templates to list and prioritize specific AI use cases and projects, but for now, confirming your North Star for AI and keeping in mind this checklist will ensure that any AI investments are intentional, coherent, and strategically meaningful. When AI is tightly coupled with business strategy, internal sponsorship and support from the rest of the leadership team and the AI committee become more natural and easy to obtain.

Let's now explore how you will need to organize your multi-level AI management approach to cover the different areas of your company, and to make sure that your vision and strategy translate into actionable items in the operational day-to-day.



## Understanding the levels of AI management

If you want to really make a difference and implement a lasting AI leadership style in your company, you need to manage the AI initiative on different levels.

Think about it as a pyramid or sandwich on three different levels: the higher level goes into long-term planning considerations, the middle level goes into tactical day-to-day planning, and the lowest level goes into technical operations. These three levels, and the way they interact with each other, are represented in *Figure 1.3*:



Figure 1.3: Levels of AI management

These levels interact and feed one another with sponsors, activities, and specific workstreams. Let's learn more about them:

- **Strategic AI management:** This level is in theory your natural habitat as an AI leader, as it defines the company's AI vision and long-term strategy, aligning AI initiatives with overall business goals, market trends, and competitive advantages. This level ensures executive sponsorship, secures funding, addresses regulatory and ethical concerns, and establishes key AI success metrics to measure business impact.

Besides elements like the AI policy and the creation of an AI committee, the strategic management level can add structures such as a **Center of Excellence (CoE)** that takes the standardized best practices and guidance and brings them to the next level.

- **Tactical AI management:** This second level translates AI strategy into actionable plans by identifying priority use cases, allocating budgets and teams, setting timelines, and managing cross-functional collaboration between technical and business units. This level ensures AI projects stay on track, evaluates risks, monitors AI performance against business KPIs, and adapts initiatives based on market or operational feedback.

This level is represented by departments such as a **Project Management Office (PMO)**, or roles like AI project managers, product managers, or technical program managers. While the scope of the strategic level is company-wide, the tactical piece focuses on specific projects and use cases. That means that the aggregated accomplishments, best practices, issues, and potential improvements can feed back into the activity of AI committees

- **Technical AI management:** The last and most technical level focuses on the hands-on development, deployment, and maintenance of AI systems, including model training, data pipeline management, infrastructure scaling, and security compliance. This level ensures AI solutions are optimized for accuracy, efficiency, and integration into existing IT systems, while continuously improving performance through monitoring, debugging, and retraining models as needed.

Technical management is tightly related to the notion of operations. Even if you have already come across terms like DevOps, MLOps, and LLMOps, in the next chapter we will go beyond the buzzwords and explore what these operational practices actually imply for your organization, along with other concepts such as the AI lifecycle (understood as the end-to-end of an AI system from ideation to maintenance and decommissioning).

Assuming you are leading the AI initiative from the top, you will need to make sure you analyze, understand, redesign, and monitor the multiple levels of management to make sure all the pieces (e.g. hierarchy, people, processes) are in place and ready to collaborate. The idea of the different levels collaborating up and down to create a continuity between levels is especially important because that continuity is what turns strategy into working systems: strategic decisions must reach the teams that build and run the technology, and operational lessons must flow back up so the strategy remains realistic. Without this two-way flow, each level optimizes in isolation, priorities diverge, and the overall AI initiative loses coherence and speed.

## Improving continuously via management systems

Often related to compliance tasks, the idea of a management system that monitors and controls the end-to-end implementation of AI at the company level is especially relevant for your AI leadership exercise, as management systems are not focused on specific projects or use cases, but on the entire AI implementation scope of the organization.

A management system is a structured framework of processes, roles, rules, and tools that an organization uses to set objectives, make decisions, execute work, monitor performance, and continuously improve. It helps define in a programmatic way how an organization turns strategy into consistent, repeatable results.

In the context of AI, management systems are part of the scope of international standards such as the **International Organization for Standardization (ISO)**. There are several standards that can help you leverage best practices for your AI implementations, including the idea of management systems. *Figure 1.4* shows some of the main ISO standards that are relevant for artificial intelligence, including ISO 42001 for AI management systems:




Figure 1.4: ISO Standards for Artificial Intelligence

Among all these standards, we will single out ISO 42001 because it is the first international, certifiable standard specifically dedicated to AI management systems, which makes it the most directly actionable one for your AI leadership exercise. Just as ISO 27001 is a standard for information security, ISO 42001 establishes a structured **AI Management System (AIMS)** that enables organizations to govern AI systems in a systematic, repeatable, and responsible manner, ensuring alignment with ethical, legal, and business objectives.

A key component of ISO 42001 is AI risk management, which addresses concerns related to bias, security, privacy, and ethical dilemmas by incorporating systematic risk identification, mitigation strategies, and continuous monitoring mechanisms such as AI auditing, drift detection, and compliance reviews.

By implementing these governance structures, ISO 42001 ensures that AI systems remain trustworthy, explainable, and resilient, ultimately supporting responsible AI deployment at scale. It serves as a foundational framework that balances innovation with ethical considerations, enabling organizations to harness AI's full potential while minimizing associated risks and maintaining stakeholder confidence in AI-driven decision-making.

Unlike individual AI projects, which focus on technical implementation, the AI management system operates on a higher management level (from strategic to tactical and technical levels), setting policies, governance structures, and risk management frameworks that guide all AI initiatives within an organization. This ensures that AI efforts are not siloed but instead contribute directly to the company's overall AI strategy, integrating seamlessly with business goals, regulatory compliance, and stakeholder expectations.

The management system provides you with a lifecycle approach to AI, covering development, deployment, monitoring, and eventual decommissioning while enforcing accountability, transparency, and risk-based governance.

The ISO standard explains the what and why, but the "how" or implementation approach depends on the company. You can build your AI management system based on the experience developed with your new AI use cases, in an incremental and iterative way by increasing the scope of the areas and products being managed by the management system. This incremental approach is already visible in the market: some of the major cloud and AI providers were among the first organizations to publicly announce ISO 42001 certifications for their AI services, and they did so largely by extending existing security and privacy management systems rather than building new ones from scratch. As a more relatable illustration, consider a mid-size insurer that starts by applying the management system to a single high-risk underwriting use case, refines its roles and controls there, and only then extends the same processes to its marketing and customer service applications.

So what should you actually do with ISO 42001? First, if your company already holds an ISO 27001 certification for information security, you are in luck: both standards share the same high-level management system structure, so you can extend existing processes, documentation habits, and audit relationships instead of starting from zero. Second, treat certification as a deliberate decision rather than a default: no regulation strictly requires ISO 42001 today, and relatively few companies worldwide have been certified so far, so early adoption is more a signal of governance maturity and a potential differentiator with enterprise customers and regulators than a compliance obligation. Third, calibrate the effort realistically: a typical journey includes a gap assessment, the implementation of the AI management system, and an external audit, and depending on your company's size and maturity, it usually spans from several months to more than a year, with the investment driven mainly by internal time and external audit work. A pragmatic path is to start with a gap assessment against the standard, align your AI policy and AI committee with its requirements, and decide later whether formal certification adds value for your context.

## Evaluating the AI management levels

Before concluding this first chapter, try to evaluate the current state of your company's three levels of AI management, based on the available information. Here is a checklist that you can use to carry out the evaluation:

### Tip

#### AI Management Evaluation Checklist

##### 1. Strategic AI Management Checklist (*Vision, governance, long-term value, and trust*)

###### ☐ a. AI Vision and Strategy

- Is there a clearly articulated AI vision aligned with corporate strategy?
- Are AI initiatives explicitly linked to business outcomes (growth, efficiency, risk reduction, innovation)?
- Have priority domains (e.g., operations, customer, product, risk) been identified for AI impact?
- Is there a multi-year AI roadmap with clear strategic milestones?

###### ☐ b. Executive Sponsorship and Decision Rights

- Is there a named executive sponsor accountable for AI at company level?
- Are AI-related decisions (investment, risk acceptance, ethics) clearly owned at executive level?
- Is AI discussed regularly in executive or board-level forums?

###### ☐ c. Governance, Ethics, and Policy

- Does the company have a formal AI policy (acceptable use, accountability, risk appetite)?
- Are ethical principles for AI (fairness, transparency, human oversight) explicitly defined?
- Are regulatory and legal requirements (e.g., data protection, AI regulation) systematically addressed?
- Is there an AI committee or governance body with cross-functional representation?

☐ d. **Organizational Enablement**

- Is there an AI Center of Excellence (CoE) or equivalent structure?
- Does the CoE define and disseminate standards, best practices, and reusable assets?
- Are AI skills and capabilities actively developed at leadership and workforce level?

☐ e. **Strategic Metrics and Value Tracking**

- Are AI success metrics defined at business level (not just technical KPIs)?
- Is there a mechanism to track portfolio-level AI value realization?
- Are underperforming or misaligned AI initiatives actively stopped or redirected?

**Add your notes here:**

---

---

---

---

## **2. Tactical AI Management Checklist** (*Execution, prioritization, and delivery control*)

☐ a. **Use Case Portfolio and Prioritization**

- Is there a structured intake and prioritization process for AI use cases?
- Are use cases evaluated based on business value, feasibility, and risk?
- Is there transparency on which AI initiatives are active, planned, or on hold?

☐ b. **Program & Project Management**

- Is AI delivery integrated into a PMO or program management structure?
- Are roles such as AI product owner / AI project manager clearly defined?
- Are dependencies between business, data, IT, and legal teams actively managed?

☐ **c. Resource and Budget Management**

- Are AI budgets explicitly allocated and tracked (not hidden in IT or innovation budgets)?
- Are teams sized and staffed according to the ambition of the AI portfolio?
- Is external vendor or partner usage governed and strategically aligned?

☐ **d. Performance Monitoring and Adaptation**

- Are AI initiatives monitored against business KPIs, not only delivery milestones?
- Is there a formal process to adapt or reprioritize AI projects based on results?
- Are lessons learned, best practices, and failures systematically captured?

☐ **e. Feedback to Strategic Level**

- Is there a structured feedback loop from projects back to AI governance bodies?
- Are recurring issues (data quality, skills gaps, tooling limits) escalated and addressed?
- Do tactical insights influence policy updates, standards, or strategy revisions?

**Add your notes here:**

---

---

---

---

### **3. Technical AI Management Checklist** *(Engineering excellence, reliability, and operational control)*

☐ **a. AI Development and Lifecycle Management**

- Is there a defined AI lifecycle from ideation to decommissioning?
- Are model development, validation, deployment, and retirement standardized?
- Is documentation maintained for models, data sources, and assumptions?



☐ **b. Data, Infrastructure and Tooling**

- Is data availability, quality, and ownership clearly defined?
- Are AI workloads supported by scalable and secure infrastructure?
- Are development and production environments clearly separated and controlled?

☐ **c. Operations (DevOps / MLOps / LLMOps)**

- Are AI systems monitored in production for performance, drift, and failures?
- Are retraining, versioning, and rollback procedures in place?
- Is there incident management for AI-related failures or unexpected behavior?

☐ **d. Security, Privacy, and Compliance**

- Are AI systems integrated into existing cybersecurity and privacy controls?
- Are access rights, audit logs, and data lineage managed?
- Are third-party models and APIs assessed for risk and compliance?

☐ **e. Technical Quality and Reliability**

- Are accuracy, robustness, and explainability measured and tracked?
- Are bias, fairness, and performance degradation actively monitored?
- Is technical debt in AI systems explicitly managed?

**Add your notes here:**

---



---



---



---

#### **4. AI Management System (Cross-Level) Checklist** (*Continuity, trust, and continuous improvement – ISO 42001-aligned*)

☐ **a. Management System Foundations**

- Is there a company-wide AI Management System (AIMS) in place or being built?

- Are roles, responsibilities, and decision rights formally defined?
- Are AI policies, standards, and procedures centrally maintained?

☐ **b. Risk Management and Assurance**

- Is there a systematic process for AI risk identification and assessment?
- Are risks related to bias, ethics, security, and compliance actively mitigated?
- Are AI systems periodically reviewed or audited?

☐ **c. Monitoring, Review, and Improvement**

- Are AI KPIs, risks, and incidents regularly reviewed at management level?
- Is there a formal continuous improvement cycle (plan–do–check–act)?
- Are findings from audits, incidents, and projects used to improve governance?

☐ **d. Lifecycle and Portfolio Oversight**

- Is the full AI portfolio visible and managed holistically?
- Are AI systems retired or decommissioned when no longer fit for purpose?
- Does the management system scale as AI adoption grows?

**Add your notes here:**

---



---



---



---

**5. Executive Expectations for AI Management Levels** (*What you should expect from the multiple levels*)

☐ **a. Expectations for Strategic AI Management (direction, legitimacy, and long-term value)**

- A clear and stable AI direction that guides investment and prioritization
- AI initiatives that are explicitly tied to business strategy
- Consistent executive messaging on AI ambition, ethics, and risk appetite

- Company-wide standards and policies that reduce fragmentation
- A single source of truth for AI governance and accountability
- Transparent reporting on AI value creation and risk exposure

☐ **b. Expectations for Tactical AI Management (translation, orchestration, and delivery discipline)**

- AI strategy translated into clear, prioritized use cases
- Realistic plans with defined scope, owners, budgets, and timelines
- Active coordination between business, IT, data, legal, and risk functions
- Early identification of execution risks and trade-offs
- Regular reporting focused on business progress, not technical detail
- Structured feedback into strategic governance bodies

☐ **c. Expectations for Technical AI Management (build, operate, and sustain AI systems responsibly)**

- Reliable, secure, and scalable AI systems in production
- Clear ownership of models, data, and pipelines
- Monitoring of performance, drift, and operational risks
- Integration of AI into existing IT and security processes
- Transparent documentation and traceability
- Continuous improvement through retraining and optimization

☐ **d. Expectations for the AI Management System (continuity, trust, and scalability)**

- A consistent way of governing AI across the organization
- Clear linkage between strategy, execution, and operations
- Risk-based governance that scales with AI adoption
- Evidence of compliance, not just good intentions
- A lifecycle view covering build, run, monitor, and retire
- Continuous improvement driven by data, audits, and experience

**Add your notes here:**

---



---



---



---

☐ **6. Executive Red Flags (Quick Self-Test):** If you answer "Yes" to any of these, you will need to note it down, pay attention, and try to deep dive and find more information from relevant stakeholders.

- ☐ "We have AI projects, but no clear AI strategy"
- ☐ "AI decisions depend on individual teams, not governance"
- ☐ "We measure model accuracy, but not business impact"
- ☐ "We don't know all AI systems currently in production"
- ☐ "Compliance and ethics are handled case by case"

**Add your notes here:**

---



---



---



---

If you're worried that you don't have all the answers mentioned in the preceding checklist yet, know that it's okay, and it's one of the reasons you picked up this book. Besides, certain parts like the management system may be totally new for you and your company. For now, just keep this checklist in mind when you revisit the AI levels of management, and try to anticipate potential gaps or improvements you may add so each level brings the expected value to the AI management exercise. A simple way to do this: for every unchecked item, note down who owns that area today and whether the gap is mainly about people, processes, or technology. That quick exercise will give you a first, informal maturity snapshot that you can revisit as you progress through the book.

## Summary

In this chapter, we learned why AI adoption must be led from the top, and why your role as an executive in that effort is not optional anymore. We started by revisiting what artificial intelligence really means today, from the classic definitions to the notion of digital intelligence, and we explored how the generative AI boom changed the level of engagement expected from leaders, including how to deal with the inevitable impostor syndrome.

We then moved to the practical part of your executive playbook: creating your first AI policy, designing and evolving a multidisciplinary AI committee, and confirming your AI North Star by aligning use cases with your corporate strategy through Porter's Generic Strategies Framework. Finally, we covered the three levels of AI management, i.e., strategic, tactical, and technical, and the value of management systems and standards such as ISO 42001 to govern AI in a continuous, risk-based manner, along with checklists to evaluate where your company stands today.

Let's now move to *Chapter 2* and explore the timeline and history of the artificial intelligence domain and industry. Get ready!

# 2

## From Turing to ChatGPT: A Brief History and Context of AI

In the first chapter, you explored a series of recommendations and templates for your AI leadership and management activities. Now it is time for you to upskill and learn about the key topic of this playbook: AI itself.

Let's be clear. While you are not required to become a true expert, having a baseline understanding of some basic and relatively advanced topics will equip you and help you make better and more informed decisions.

Imagine situations in which you will get two competitive requests for resources, and you need to make sense of the nature of both AI projects. If you don't understand the basic technology components, you won't be able to analyze their feasibility and measure the potential cost of opportunity. Or situations in which some teams or clients may propose the use of very new technologies, in which case you may either face an innovation dilemma (waiting until the technology is a bit more mature) or unjustified optimism.

This is why we will incrementally build your AI literacy between *Chapters 2* and *4*, so you can lead your AI discussions with confidence and the ability to analyze potential options, gaps, and alternatives. This chapter focuses on the historical perspective of AI innovations and the considerations you need to keep in mind in this era of generative AI, and *Chapters 3* and *4* will deep dive and explain complex technology topics in simple (at least as simple as possible) words.


Specifically, in this chapter, we will cover the following topics:

- Historical perspective
- Key milestones from an AI adoption perspective
- How generative AI changes your game

- Comparison with the printing press
- Lessons from past hype cycles
- Geopolitics and power in the context of AI

## Historical perspective

In general terms, you can see the history of artificial intelligence as a series of decades that include relevant trends and milestones, ups and downs, and an increasing level of technical capabilities that were accompanied by higher levels of industry adoption. This evolution is summarized in the timeline in *Figure 2.1*:



*Figure 2.1: AI Timeline*

All together created the base of what AI means today, but each phase will bring relevant information points that can help you define your AI leadership style. Let's look at each of these phases in detail.

## Pre-1950: Imagining intelligent machines before AI existed

Long before the term AI entered the scientific vocabulary, humanity was already fascinated by the idea of non-human intelligence. This fascination emerged simultaneously in philosophy, literature, engineering, and early computational research. Ancient myths of automatons, mechanical dolls, and early calculating machines all reflected a persistent desire to replicate human cognition and agency.

In the early 20th century, science fiction played a crucial role in shaping the collective imagination around intelligent machines. Works such as Karel Čapek's *R.U.R.* (1920) presented artificial beings endowed with emotions, reasoning, and moral dilemmas, while Isaac Asimov's *I, Robot* (1950) explored ethical frameworks governing intelligent machines through his famous *Three Laws of Robotics*, which first appeared in *Runaround* (1942). These narratives established a humanoid and anthropomorphic vision of intelligence, one that continues to influence public expectations of AI today.

At the same time, foundational academic work was laying the groundwork for future AI research. Advances in formal logic, mathematics, and early computing introduced the idea that machines could, in principle, exhibit intelligent behavior. In 1950, Alan Turing proposed the "Imitation Game," later known as the **Turing Test**. In Turing's paper, the imitation game starts as a party game in which an interrogator must tell a man from a woman through written notes; Turing then asks what happens when a machine takes the man's place. The commonly cited version states that the imitation game involved three participants: a human interrogator holding text-based conversations with both a machine and another human, trying to determine which was which. The goal was never to establish whether the machine truly thinks, but whether its behavior could become indistinguishable from that of a human in conversation. This is a subtle but important distinction, and one that is still widely misunderstood today.

It is worth pausing on one implication of this pre-history: many of our ethical assumptions about AI were inherited long before the technology existed. Fiction taught society to expect artificial minds that rebel, deceive, or demand rights, and frameworks such as Asimov's laws framed safety as a problem of constraining humanoid agents. As a result, public debates about AI often import fears and expectations that do not match how today's systems actually work.

As an AI leader, expect your stakeholders to arrive with these inherited assumptions, and treat correcting them (through education, transparent communication, and realistic demonstrations) as part of your job rather than a distraction from it.



**Tip****Impact of AI history on leadership (1 of 8)****Questions for self-reflection:**

- How much of today's public expectation of AI is shaped by fiction rather than reality?
- Are we still unconsciously building toward a humanoid or anthropomorphic ideal of intelligence?
- How do narratives and metaphors shape trust, fear, and adoption more than technical facts?
- What ethical assumptions did we inherit before the technology even existed?
- As an AI leader, how to overcome this when managing teams, clients, end users, and other stakeholders?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- ☐ Managing expectations vs. reality is a constant task; media, users, and investors often expect "human-like" intelligence.
- ☐ Understanding ethical discussions that usually precede technical ones because society formed opinions long before the AI capability existed.
- ☐ Spending significant time on storytelling and framing: explaining what AI is and is not. Balancing inspiration and caution, as stakeholders arrive with expectations formed by decades of fiction and media narratives.

How are you planning to approach this? What are the implications for your own leadership style and the company's AI strategy?

Add your notes here:

---

---

---

---

## The 1950s: The birth of AI as a discipline

The 1950s marked the formal birth of AI as an academic field. During this decade, rapid advancements in digital computing, combined with post-war investment in scientific research, created fertile ground for ambitious ideas about machine intelligence.

The defining moment came in 1956 with the *Dartmouth Summer Research Project on Artificial Intelligence*. Organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, this workshop was explicitly designed to establish a new, neutral term to unify research that had previously been scattered across concepts such as "cybernetics," "automata," and "thinking machines." Refer to the full proposal document at [https://parhamdata.com/Dartmouth\\_1955\\_Proposal\\_for\\_AI\\_Research.pdf](https://parhamdata.com/Dartmouth_1955_Proposal_for_AI_Research.pdf). The first page of the original proposal is shown in *Figure 2.2*:

A PROPOSAL FOR THE  
DARTMOUTH SUMMER RESEARCH PROJECT  
ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College  
M. L. Minsky, Harvard University  
N. Rochester, I. B. M. Corporation  
C. E. Shannon, Bell Telephone Laboratories

*Figure 2.2: A Proposal for the Dartmouth Summer Research Project on AI*

The proposal document for the conference was remarkably bold. It suggested that every aspect of learning, reasoning, and intelligence could, in principle, be described precisely enough for a machine to simulate it. Even by today's standards, the document reads as visionary, anticipating ideas such as symbolic reasoning, language understanding, and self-improving systems. The conference brought together leading thinkers and helped legitimize AI as a serious scientific endeavor, setting a research agenda that would dominate the following decades.

**Tip****Impact of AI history on leadership (2 of 8)****Questions for self-reflection:**

- Are today's AI leaders repeating the same pattern of bold claims made in 1956?
- How do early definitions constrain future innovation?
- Is access to AI knowledge more accessible now?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Executives should spend time choosing words carefully: "intelligence," "reasoning," "agents," "autonomy."
- Strategy and business decks may echo Dartmouth-level ambition, but internal roadmaps stay more conservative.
- Leaders must balance visionary messaging with credibility to avoid future backlash.
- Naming new products or capabilities is treated as a long-term strategic decision, not marketing fluff.
- Revisiting definitions periodically, as the meaning of terms such as "agent" or "reasoning" evolves faster than corporate glossaries.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## The 1950s to the early 1970s: Optimism, funding, and early limitations

Following the Dartmouth Conference, AI entered a period of intense optimism and generous funding, particularly in the United States. Governments and corporations alike believed intelligent machines could soon automate complex tasks, enhance military capabilities, and transform industrial productivity.

Research during this era focused largely on symbolic AI and rule-based systems. These systems relied on explicitly coded rules and logic to solve narrowly defined problems, such as playing games, solving algebraic equations, or proving theorems. While impressive in controlled environments, these systems were brittle and failed when faced with ambiguity, noise, or real-world complexity.

A pivotal moment came in 1969 with the publication of *Perceptrons* by Marvin Minsky and Seymour Papert. The book rigorously analyzed the limitations of single-layer perceptrons, an early neural network model, showing that they could not represent certain basic functions. Although technically accurate, the work had an outsized impact: it reinforced a shift that was already underway, as funding and attention were consolidating around symbolic approaches, and it contributed to keeping connectionist research marginal for more than a decade.

### Tip

#### Impact of AI history on leadership (3 of 8)

##### Questions for self-reflection:

- Like the early rule-based systems of this era, are we building demos that impress rather than systems that scale?
- Which parts of our AI stack are as brittle as those early symbolic systems, and overly dependent on assumptions?

- The perceptron episode froze an entire research field. How do we prevent narrow successes or failures from being mistaken for general capability?
- Are we overfunding AI initiatives based on excitement, as governments and corporations did in the 1960s?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Managing constant tension between research breakthroughs and innovation, and production reliability.
- Pushing teams to move beyond "cool demos" toward robust systems.
- Investing upfront in evaluation, monitoring, and failure analysis.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## Up to 1980: The first AI winter

As ambitious promises failed to materialize into practical, scalable systems, skepticism toward AI grew. This period saw the emergence of what later became known as the first AI winter. This phase was characterized by reduced funding, declining public interest, and critical reassessment of the field.

A major turning point was the Lighthill Report in 1973, commissioned by the British government. The report sharply criticized AI research, arguing that progress had been minimal relative to investment and that most successes were limited to toy problems with no real-world

applicability. As a direct consequence, the UK significantly cut funding for AI research, and similar sentiments spread internationally.

The combination of unmet expectations, technical limitations, and political scrutiny led to a prolonged slowdown. AI did not disappear, but it retreated into smaller academic niches, awaiting new ideas and enabling technologies.

**Tip****Impact of AI history on leadership (4 of 8)****Questions for self-reflection:**

- What signals, similar to those that preceded the first AI winter, would indicate that we're entering another hype correction?
- Which of our investments would survive a funding or sentiment downturn like the one that followed the Lighthill Report?
- How do we communicate setbacks without killing momentum, avoiding the unmet-expectations trap of the 1970s?
- What happens if regulators or governments lose confidence in AI progress? What is your plan?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Actively managing expectation risk with boards and investors.
- Focusing on defensible value creation, not just future promises.
- Planning for downside scenarios: budget tightening, slower adoption, regulatory friction.
- Focusing on internal morale and retention during moments of public skepticism.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## The 1980s: Expert systems, revival, and the second winter

The early 1980s brought a renewed surge of optimism driven by practical, commercially viable AI applications. Chief among these were expert systems, which are rule-based programs designed to replicate the decision-making abilities of human experts in narrow domains such as medical diagnosis, geology, or industrial troubleshooting. These systems found real-world adoption and attracted significant corporate investment.

At the same time, Japan launched the Fifth Generation Computer Systems Project, an ambitious national initiative aimed at creating AI-native computers capable of logic programming and intelligent reasoning. This project further intensified global interest and competition in AI research.

Crucially, the mid-1980s also witnessed the rediscovery of neural networks. In 1986, David Rumelhart, Geoffrey Hinton, and Ronald Williams demonstrated the practical effectiveness of backpropagation for training multilayer neural networks (the underlying technique had been independently derived by earlier researchers), overcoming some of the limitations of single-layer perceptrons that Minsky and Papert had highlighted in their 1969 book. Despite these breakthroughs, expert systems proved expensive to maintain and difficult to scale, and the promised hardware revolutions failed to deliver.

Unfortunately, by the late 1980s, commercial enthusiasm collapsed once again, leading to the second AI winter, generally dated between 1987 and 1993.

### Tip

#### Impact of AI history on leadership (5 of 8)

##### Questions for self-reflection:

- Are our systems, like the expert systems of the 1980s, too expensive to maintain at scale?
- What technical debt are we creating today that will break tomorrow, as the hand-crafted rule bases of expert systems eventually did?

- Are we mistaking rule complexity or static workflows for intelligence?
- How resilient is our business model if the technology plateaus, as it did before the second AI winter?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Monitoring maintenance cost, latency, and scalability.
- Pushing toward modular architectures and automation.
- Influence product decisions by total cost of ownership.
- Avoid over-specialized, fragile teams.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## Between the 1990s and 2006: From hype to real-world impact

The 1990s marked a more mature phase for AI, characterized by steady, application-driven progress rather than grand promises. Machine learning shifted toward statistically grounded methods, including Support Vector Machines, decision trees, and ensemble models such as Random Forests. These techniques proved robust, interpretable, and effective across industries like finance, insurance, and manufacturing.



Simultaneously, the rise of the internet transformed data availability. The explosion of digital content and user-generated data laid the foundation for large-scale learning systems and the emergence of what would later be called big data.

This era also saw several landmark achievements. In 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov, a milestone achieved by a sophisticated special-purpose system that combined custom chips capable of evaluating hundreds of millions of chess positions per second with opening books, endgame databases, and evaluation functions tuned with input from grandmasters. There is also a leadership lesson hidden in this story: after the public relations triumph, IBM retired Deep Blue without turning it into a sustained line of business, a useful reminder of the difference between a one-off innovation showcase and a durable engine of value creation. In parallel, advances in natural language processing and computer vision progressed steadily, with convolutional neural networks gaining traction for image recognition tasks.

The culmination of this period arrived in 2006, when Geoffrey Hinton and his collaborators introduced deep belief networks, using unsupervised, layer-by-layer pre-training based on restricted Boltzmann machines to overcome the vanishing gradient problem that had long prevented deep networks from being trained effectively. This signaled a renewed and scalable approach to deep neural networks. This moment is often cited as the beginning of the modern AI revolution. Notice the pattern: a problem the field considered closed ("neural networks don't scale") was reopened by a fundamentally different approach, something worth remembering whenever your organization labels a problem as impossible.

**Tip****Impact of AI history on leadership (6 of 8)****Questions for self-reflection:**

- Following the pragmatic spirit of the 1990s, where does AI create measurable ROI today?
- Which of our problems are better solved with the robust statistical methods of that era rather than with the latest models?
- How do we integrate AI quietly into workflows, as companies did back then, without overselling it?
- Like the internet-era pioneers of big data, how do we turn data into a long-term strategic moat?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Focusing on AI as infrastructure, not spectacle.
- Collaborating with data, legal, and operations teams.
- Defining KPIs that revolve around efficiency, accuracy, and cost reduction.
- Embedding AI in products without it being the headline feature.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## Between 2006 and the 2010s: The deep learning revolution

From 2006 onward, **deep learning (DL)** became the dominant paradigm in AI research and applications. Three factors converged to enable this shift: vastly increased computational power (notably GPUs), massive labeled datasets, and new approaches for training deeper neural networks. Individually, none of these factors was enough to transform the field. Together, they changed the trajectory of AI.

The defining moment came in 2012 with AlexNet, a deep neural network developed by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. In the ImageNet competition, AlexNet reduced the top-five classification error rate from roughly 26 percent to about 16 percent, a performance leap that no incremental improvement had achieved in previous years. This success validated deep neural networks at scale and triggered widespread adoption across academia and industry. This was the moment that DL stopped looking like an academic curiosity and started looking like a viable foundation for future products and services. Rapid advances followed in speech recognition, computer vision, and natural language processing.

This achievement illustrates how innovation compounds over time. Alex Krizhevsky, at the time a doctoral student in Geoffrey Hinton's group, built on years of foundational research into DL, including the 2006 deep belief network breakthrough we just discussed. What appeared to be a sudden breakthrough in 2012 was actually the result of a much longer journey of research, experimentation and technological progress. This serves as an important reminder that many transformative innovations emerge not from a single discovery, but from accumulation of advances that eventually reach a tipping point.

The leadership lesson from AlexNet is not about image recognition. It is about convergence. Neural networks had existed for decades prior, yet they only became transformative when advances in data availability, computing power and training techniques matured simultaneously. Many technologies follow the same pattern. Leaders often search for a single breakthrough moment, when in reality the most disruptive innovations emerge when several independent trends reinforce one another. Understanding these convergence points is often more valuable than tracking any individual technology in isolation.

In 2016, AlphaGo's victory over world champion Lee Sedol showcased the power of combining deep reinforcement learning with neural networks. However, perhaps the most transformative milestone arrived in 2017 with the publication of *Attention Is All You Need*. The introduction of the Transformer architecture and self-attention mechanisms revolutionized sequence modeling and laid the foundation for **large language models (LLMs)**.

**Tip****Impact of AI history on leadership (7 of 8)****Questions for self-reflection:**

- Given that DL rewards scale, do we have the data, compute, and talent to compete at scale?
- As models grow larger, when does performance improvement justify exponential cost?
- Are we focusing on individual models, or on platforms that can absorb the next AlphaGo-style breakthrough?
- Considering that GPUs enabled the DL revolution, how dependent are we on hardware and cloud providers?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Dedicating time to compute strategy and partnerships.
- Hiring and retaining top AI talent becomes existential.
- Navigating tradeoffs between open research and competitive advantage.
- Planning long-term infrastructure details.

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## The 2020s: AI goes mainstream

With a mature technological foundation in place, the early 2020s ushered in an era of explosive growth and mainstream adoption. AI moved beyond specialized domains to become a ubiquitous presence across society, influencing education, healthcare, science, creative industries, and everyday productivity.

The release of OpenAI's ChatGPT in 2022 marked a defining moment, bringing conversational AI into the public consciousness at an unprecedented scale. This catalyzed a wave of competition and innovation, with rapid advancements from organizations such as Anthropic, Google, Meta, Mistral, and others. New generations of language, vision, speech, and multimodal models emerged at an extraordinary pace.

It is worth distinguishing here between a technology breakthrough and a market breakthrough. ChatGPT was primarily the latter: the underlying research had been maturing for years, with GPT-3 already showing the potential of large-scale language models in 2020, and GPT-4 raising the technical bar again in March 2023. The two types of breakthroughs demand different strategic responses from you as a leader. Technology breakthroughs change what is possible and call for exploration and evaluation, while market breakthroughs change what customers,

employees, and boards expect, and call for positioning, governance, and adoption decisions at a much faster pace.

This period has been characterized by massive investment, intense competition, and rising expectations, alongside growing debates about regulation, ethics, and long-term societal impact. As AI continues to evolve, the 2020s stand as a transformative decade in which AI has shifted from a specialized technology to a foundational layer of the entire company ecosystem.

**Tip****Impact of AI history on leadership (8 of 8)****Questions for self-reflection:**

- How do we deploy fast without breaking trust?
- Where do we draw the line between capability and responsibility?
- How do we decide when everyone can have access to the AI capabilities?
- What is our stance on regulation, safety, and alignment?
- How do we scale adoption without overwhelming users?

**Add your notes here:**

---

---

---

---

**Which of the following factors has an impact on your AI leadership style?**

- Spending as much time on policy, ethics, and public discourse as on tech.
- Accepting crisis management as routine (hallucinations, misuse, public backlash).
- Adding pricing, distribution, and partnerships to AI strategy discussions.
- Finding your balance - keeping regular pace vs. operating in permanent beta: shipping, learning, correcting at speed.
- Defining AI leadership as cross-functional by default (tech, legal, comms, policy, infra).

How are you planning to approach this? What are the implications for your own leadership style, and the company's AI strategy?

**Add your notes here:**

---

---

---

---

## Reading the pattern: What AI history teaches your business

You can find every date and milestone mentioned above in an encyclopedia or by asking a chatbot. What matters for you as a leader is the pattern that emerges when you connect them, and there are three lessons I want you to retain from this journey.

First, enterprise AI succeeds or fails for business reasons more often than for technical ones. Expert systems in the 1980s did not collapse because their rules were wrong; they collapsed because maintenance costs grew faster than the value they delivered, because the knowledge they encoded was locked inside a few scarce specialists, and because they could not adapt when the business around them changed. Replace "rules" with "prompts and integrations," and you have a fairly accurate preview of how today's LLM deployments fail: impressive pilots that nobody budgets to maintain, evaluate, or own. When you review an AI proposal, ask about the operating model behind the demo.

Second, AI winters were never caused by the technology alone, but by the gap between promises and deployed value. Funding follows narrative, and narrative overshoots in both directions. The organizations that crossed previous winters comfortably were those holding measurable, defensible value when sentiment turned. Build your portfolio so that at least part of it would survive a change in narrative.

Third, breakthroughs come from convergence and compounding, not from isolated genius. Backpropagation needed decades, cheap GPUs, and internet-scale data before it changed anything; the 2006 pre-training ideas matured into AlexNet, the Transformer architecture, and eventually ChatGPT. This means the most useful thing you can track is not any single technology, but the convergence of enablers: data availability, compute costs, and distribution channels. When several of those curves bend at once, a tipping point is near, and that is when timing decisions matter most.

All these stages and decades brought new collective learnings that are very useful nowadays, especially when trying to avoid past mistakes and focusing on pragmatic yet meaningful AI innovation for your company.

The rest of the chapter will focus on specific topics you should keep in mind, as they will be part of your external and internal context during the leadership and adoption exercise. Let's start by looking at which of these historical milestones actually changed the way companies adopt AI, and why the gap between research and commercial impact keeps shrinking.

## **Key milestones from an AI adoption perspective**

All these eras and accomplishments are relevant to understanding what AI is today, but it doesn't mean that all the research milestones and innovations have equal impact on your day-to-day life. There are plenty of scientific papers and incremental improvements that take years to reach to the commercial sphere.

However, there are other innovations that have clearly impacted the way companies and end customers leverage AI. And the time delta between the research work and commercial release of certain technologies has drastically reduced in the last 10 years.

For example, the machine learning models that you will learn about in the next chapter took their time to deeply impact the industry. The 1990s brought a series of capabilities that companies from different industries (e.g., finance, insurance) started to leverage at a sustainable pace. If you think about it, we are still seeing new cycles of adoption and use cases for those industries and technologies. Longer period, more mature adoption patterns.

On the other hand, technologies like DL or LLMs took just a few years to reach the masses. Applications like the facial recognition features on your phone, the picture classification in your cloud folder, or the omnipresent LLM-enabled chats like ChatGPT, Copilot, or Gemini went from the research labs to the end products and users in very little time. And the pace is only accelerating now that technology providers and the industry in general are rushing to produce the best models, applications, and use cases.

This acceleration from research lab to mass adoption is precisely what makes the current wave different from everything you just read about. Now that you have seen how quickly research becomes product, let's explore what generative AI specifically changes for you as a leader.

## **How generative AI changes your game**

There are different reasons to believe that the generative AI wave has some similarities with other past innovations (e.g., buzz, huge investments), but the reality is that the adoption game (and therefore your leadership approach) is totally different.

Think about this: the way a company or even an individual leverages AI is totally different nowadays. Before, it was mostly about technical folks in their very technical environments. Today, AI adoption is unleashed, and you can find it everywhere, including your organization.

Here are some of the key aspects you want to keep in mind before you explore the technical foundations of generative AI in *Chapter 4*:

- **Lower adoption barrier because language is the new interface.** For the first time in the history of computing, the primary interface to advanced technology is natural language itself. This fundamentally changes the adoption curve of AI because users no longer need to learn programming languages, complex tools, or technical abstractions to extract value. Asking, instructing, refining, and iterating through plain language collapses the distance between intent and execution. As a result, AI adoption is no longer gated by technical literacy but by curiosity, clarity of thought, and organizational permission. That said, this insight should not be read as absolute: natural language interfaces are non-deterministic (the same prompt can produce different results), harder to audit than traditional software, and can create a false sense of precision. Anyone can prompt a model, but not everyone can prompt it well, so the skill gap has not disappeared; it has shifted from writing code to thinking clearly and specifying precisely. For leaders, this means AI spreads horizontally across the organization at a speed traditional change management models were never designed to handle, forcing a shift from controlled rollouts to guided enablement.
- **One model and deployment, multiple uses and simplified adoption.** Unlike previous enterprise technologies that required specialized systems for each function, modern foundation models can support a wide range of use cases from a single deployment. The same model can assist with coding, marketing copy, legal analysis, customer support, and strategic planning, dramatically reducing integration complexity. This generality lowers both technical and organizational friction: teams do not need to wait for bespoke AI solutions to be built for them. This point supports a top-down leadership approach to technical operations, as there will likely be a platform and model convergence between most of your AI projects.
- **AI reaches non-technical departments.** Generative AI breaks the historical pattern where advanced technology adoption was driven almost exclusively by engineering or data teams. Today, marketing, sales, HR, finance, legal, and operations can independently integrate AI into their daily workflows, often without central approval. This democratization accelerates productivity but also fragments visibility and control. From a leadership perspective, AI is no longer an IT initiative, but an organizational capability that reshapes how work is done, decisions are made, and performance is



evaluated. And of course you must think in terms of cultural adoption, cross-functional standards, and shared principles rather than isolated technical roadmaps.

- **The talent equation has changed.** The competition for AI talent is one of the most intense the labor market has seen, with aggressive compensation packages, acquisitions designed mainly to absorb teams, and a widening gap between AI-literate and AI-illiterate organizations. At the same time, generative AI redefines what "AI talent" means: you need fewer people who can build models from scratch and more people who can apply, evaluate, integrate, and govern them. This has two implications for you. First, retention and upskilling become as important as hiring, because your existing workforce is your largest pool of potential AI practitioners. Second, AI literacy at the leadership level is itself a talent issue, and the upskilling journey you are undertaking with this book is part of the answer.
- **Multimodality is creating new opportunities; data is king.** Modern AI systems are no longer limited to text; they can process and generate images, audio, video, code, and structured data within a single workflow. This multimodality enables entirely new classes of products, experiences, and internal tools, collapsing previously separate creative and analytical pipelines. For businesses, this means faster ideation, richer customer interactions, and new ways to automate complex, multi-step processes. For leaders, it signals a shift from incremental optimization to rethinking end-to-end value chains, as AI becomes capable of operating across formats that once required multiple specialized teams or vendors.
- **Any company, small or big, can create something in hours.** The combination of pre-trained models, APIs, and cloud-based tooling has drastically reduced the time and resources required to build AI-powered products or internal tools. What once took months of research, infrastructure setup, and hiring can now be prototyped in hours. This levels the competitive playing field, allowing startups and small teams to move at speeds previously reserved for large tech companies. This means competitive advantage is increasingly temporary and execution speed becomes a strategic weapon. The barrier is no longer access to technology, but clarity of purpose and the ability to rapidly test, learn, and iterate.
- **Experimentation cost is radically lower, and innovation is more accessible.** The cost of trying, failing, and iterating with AI has dropped by orders of magnitude. Teams can now explore ideas, validate assumptions, and test workflows with minimal upfront investment. This encourages a culture of experimentation, but it also increases noise: not every experiment creates real value. So the challenge is no longer how to enable experimentation, but how to channel it toward meaningful outcomes. Governance shifts from approval-based control to outcome-based evaluation, where learning

velocity and impact matter more than polished business cases. Or even creating internal ventures within your organization based on newly developed technologies and use cases.

- **The market is fully adapted to OPEX, while providers focus on heavy CAPEX.** Except if you are building your own infrastructure capabilities, AI is consumed as an **operational expenditure (OPEX)**: usage-based pricing, subscriptions, and flexible scaling are now the norm. Behind the scenes, however, AI providers are making enormous **capital expenditure (CAPEX)** investments in data centers, specialized hardware, and energy infrastructure. This asymmetry has important strategic implications. You need to understand that while AI feels lightweight and inexpensive to adopt, it is underpinned by massive fixed costs that influence pricing, availability, and long-term vendor dynamics. Also, be aware that OPEX-based pricing can create a false sense of cost predictability, as usage-based AI costs can scale non-linearly as adoption grows, depending on factors such as context window usage, model size, and concurrency levels. More than one company has been surprised by its inference bill when moving from a small pilot to a company-wide deployment. Strategic dependency on providers, cost predictability, and negotiating power become core executive concerns, not procurement details. But don't worry, we will talk about handling vendors and tools in *Chapter 8*.
- **Time to AI is shorter, just minutes.** Accessing state-of-the-art AI capabilities no longer requires long procurement cycles, infrastructure provisioning, or specialized deployment teams. In many cases, an API key and a few lines of code (or even a no-code interface) are enough to bring powerful AI into production workflows. This compression of time-to-value accelerates innovation but also bypasses traditional organizational safeguards. You and the rest of the leadership team must adapt to a reality where AI can enter the company faster than policies can be written, making proactive guidance, clear principles, and rapid feedback loops more effective than slow, centralized approval processes.
- **More value means more risk.** As AI becomes more capable and deeply embedded in decision-making, the potential upside grows alongside the potential downside. Errors, biases, hallucinations, or misuse can now propagate at scale, affecting customers, employees, and public perception simultaneously. The same systems that generate efficiency and creativity can also introduce legal, ethical, and reputational risks. For leadership, this requires a reframing of risk management: AI risk is not a future concern or a technical edge case, but a recurrent, present, systemic issue that must be addressed continuously.

- **AI is everywhere and it impacts everything.** Generative AI is no longer confined to a single product line or operational function, It touches reputation, brand voice, innovation pipelines, pricing models, financial forecasting, and even organizational identity. Decisions about AI adoption implicitly become decisions about how the company thinks, communicates, and competes. As you already know, this means that AI strategy cannot be delegated or isolated, and it must be integrated into core leadership responsibilities. The question is no longer whether AI will impact the business, but whether leadership is intentionally shaping that impact or reacting to it after the fact.
- **AI brings new investor funding, but make sure you use that power wisely.** The current generative AI wave has unlocked unprecedented access to capital, visibility, and strategic partnerships for companies that can credibly position themselves within the AI narrative. Investors, boards, and sponsors are actively seeking exposure to AI-driven growth, often rewarding ambition, speed, and vision. However, this attention comes with heightened expectations and long-term obligations. Overpromising, vague positioning, or superficial AI adoption may generate short-term enthusiasm (if any) but can quickly erode trust when results fail to materialize. Keep this in mind while you leverage the investor (and any other stakeholder) excitement in your favor.

Summarizing, the emergence of generative models represents a tipping point never seen before. And that's good in terms of AI adoption, setting a great and nice scenario for you to bring the most value in terms of AI leaders.

Like it or not, but all these reasons are making generative AI the superstar of all the shows. The level of attention is unprecedented, but not its impact as a technology. One of the recurrent comparisons is between generative AI (and AI in general) and one of the original mainstream innovations: the printing press. Let's look at this aspect next.

## Comparison with the printing press


Created by Johannes Gutenberg around the 1440s, the printing press was a total disruption because it allowed the reproduction of text and images mechanically and in large quantities, rather than being copied by hand, which was how most books in Europe were still produced at the time (block printing existed, and movable type had already been used in China and Korea centuries earlier).

That invention transformed how knowledge was created, shared, and preserved, and made information accessible at a scale humanity had never experienced before. But it looks like history repeats itself with generative AI, as it enables access to information at scale, with lower barriers to access all types of knowledge. All of this is accompanied by the huge impact on the global

economy and the highly illustrative pace of innovations humanity has seen since the nineteenth century.

To be fair, this comparison has become so common in AI writing that it risks turning into a cliché, so it is worth acknowledging where the analogy breaks down. The printing press democratized the distribution of existing knowledge, while generative AI creates new content, including content that may be wrong or entirely fabricated. A printed book also carried a clear signal of authorship and accountability, whereas AI-generated content often has an ambiguous origin and an uncertain relationship with the truth. Keeping both the similarities and the differences in mind will make this analogy far more credible when you use it with boards, teams, or clients who have already heard it a dozen times.

To put this discussion in perspective, the following graph places AI within the broader sequence of technological revolutions, from the steam engine to the internet. The idea I want you to retain from it is the pace of change: each successive revolution has reached mass adoption faster than the previous one, and AI is following, and accelerating, that same pattern, as shown in *Figure 2.3*:



*Figure 2.3: Technological revolutions; Image adapted from Zanden – Maddison Project Database 2023 – with major processing by Our World in Data*

This accelerating pace also explains why waves of excitement and of disappointment keep repeating with each new technology. Because we have seen this movie before, history offers

concrete lessons about how to handle technology hype, and that is exactly what we will explore next.

## Lessons from past hype cycles

Hype around AI is not a new phenomenon. Similar to the recurring cycles of "AI summers" and "AI winters," companies across industries have repeatedly faced moments of intense excitement driven by emerging technologies such as deep learning, blockchain, the metaverse, NFTs, and many other trends that attracted disproportionate attention from investors, technology providers, and early adopters. While each wave differed in technical substance, they shared a common pattern: inflated expectations, rapid experimentation, uneven value creation, and eventual normalization. Understanding these cycles is essential for leaders who want to adopt generative AI strategically rather than reactively.

Here are a few lessons from hype cycles for strategic AI adoption, grouped into three blocks (how you make adoption decisions, how you read external signals and pressures, and how you turn hype into an advantage), that you should keep in mind:

- **The innovation dilemma:** Leaders like you are often forced to choose between moving early to capture advantage or waiting to reduce uncertainty and risk. Hype intensifies this dilemma by compressing decision timelines and amplifying fear of missing out. The lesson from past cycles is that early adoption only creates value when paired with clear strategic intent. Moving fast without direction often results in scattered initiatives, while moving too slowly risks irrelevance. The real challenge is not speed itself, but alignment between innovation efforts and long-term business objectives.
- **Creating versus leveraging to mitigate risks:** During hype cycles, organizations frequently feel pressured to build proprietary solutions to signal innovation leadership. However, history shows that leveraging existing platforms, ecosystems, and providers often reduces technical, financial, and operational risk, especially in the early stages of a technology's maturation. Strategic advantage rarely comes from building everything from scratch, but from knowing when to create, when to partner, and when to adopt. Leaders must resist the assumption that ownership equals control and instead focus on flexibility and learning velocity.
- **Going all in without a backup plan:** Hype-driven decisions tend to favor bold, irreversible commitments: large investments, sweeping reorganizations, and public positioning that leaves little room for adjustment. When the narrative shifts, these decisions become liabilities. Past technology cycles teach that optionality is a strategic asset. Successful adopters design exit paths, modular architectures, and fallback

scenarios from the beginning. This does not signal lack of conviction; it reflects disciplined leadership in uncertain environments.

The next group of lessons deals with how you read and respond to external signals and pressures, from markets and investors to industry benchmarks:

- **Following everything the market and investors say:** Markets and investors are powerful signals, but they are not infallible guides to timing or relevance. Many hyped technologies were overfunded long before product–market fit existed, while others delivered value quietly outside the spotlight. Leaders who blindly follow external pressure risk prioritizing optics over substance. The lesson is not to ignore the market, but to contextualize it by balancing external narratives with internal data, customer needs, and organizational readiness.
- **Not adapting innovation timing for hyped technologies to your own context:** One of the most common mistakes during hype cycles is assuming that adoption timelines are universal. In reality, the right moment to adopt depends on industry dynamics, regulatory exposure, talent availability, and organizational maturity. Technologies do not create value in isolation; they interact with existing processes, culture, and constraints. Strategic leaders adapt timing to their specific context rather than racing toward an abstract industry benchmark.
- **Adopting because "we have to do so":** Perhaps the most dangerous driver of adoption is obligation without purpose. When organizations adopt AI simply to signal modernity or avoid appearing outdated, initiatives lack ownership, metrics, and accountability. History shows that adoption driven by fear rarely produces sustainable outcomes. The most successful transformations begin with a clear articulation of why the technology matters, where it creates value, and how success will be measured.
- **Extremes are never good, avoid both excessive pessimism and blind optimism:** Hype cycles tend to polarize decision-making, pushing leaders toward either uncritical enthusiasm or defensive skepticism. Both extremes are equally damaging. Over-optimism leads to inflated expectations, rushed deployments, and fragile commitments, while excessive pessimism results in missed opportunities and strategic stagnation. Past innovation waves show that sustainable advantage comes from calibrated confidence: acknowledging real limitations without dismissing genuine potential. Effective AI leadership requires accepting that the technology is powerful and imperfect, and making decisions grounded in evidence, experimentation, and continuous learning rather than emotional reactions to the surrounding noise.

Finally, this last set of lessons focuses on turning hype and market signals into a practical advantage:

- **Leverage market signals such as hype cycles and analyst frameworks:** Market signals like the Gartner Hype Cycle, industry reports, and analyst forecasts are valuable tools for sense-making during periods of rapid technological change. They provide a shared language to understand maturity, expectations, and risk across stakeholders. However, their true value lies not in prediction, but in perspective. You can use these frameworks as contextual inputs (not decision engines, as they are not perfect) to challenge assumptions, calibrate timing, and communicate tradeoffs internally. When interpreted critically, market signals help separate structural trends from short-term noise and support more disciplined, evidence-based AI adoption strategies.
- **Read the context when dealing with providers and technology partners:** Technology providers operate under their own constraints, incentives, and pressures, including aggressive growth targets and fundraising timelines, competitive positioning, and sales quotas. During hype cycles, these pressures intensify, often shaping how capabilities are presented and roadmaps are framed. Strategic leaders look beyond product promises to understand the provider's context: financial sustainability, dependency on continued hype, technical maturity, and long-term alignment. This perspective enables more balanced negotiations, realistic expectations, and partnerships built on transparency rather than momentum.
- **Hype and low-cost experimentation can be allies when used intentionally:** Especially in the context of AI and generative AI, while hype is often framed as a distraction, it can also serve as a powerful catalyst when paired with disciplined experimentation. Heightened attention lowers organizational resistance, unlocks budget, attracts talent, and creates permission to test ideas that might otherwise be blocked. At the same time, the dramatically reduced cost of experimentation allows your teams to explore new AI use cases, validate assumptions, and learn quickly with minimal downside.

Remember, you will never be free of risk while dealing with hyped technologies, but it doesn't mean all of them are not worth experimentation and exploration. As an AI leader, you are uniquely positioned to set the temperature and guide the teams from the top, so they understand when to analyze, hesitate, validate, or accelerate to make the most of those technologies.

## Geopolitics and power in the context of AI

It is not every day that the geopolitical context becomes so central to technological innovation. In the case of AI, the industry has evolved far beyond discussions of models, algorithms, or

performance benchmarks. AI now sits at the intersection of economics, national security, diplomacy, and global power dynamics, signaling its role as a strategic asset rather than a purely commercial technology.

The key geopolitical dimensions shaping the AI landscape can be grouped into four areas: governance and AI diplomacy, competition between regions and states, resources and infrastructure, and ethical and societal aspects. Let's start with governance and AI diplomacy, which covers how governments, institutions, and technology leaders negotiate the rules of the game. The dimensions in this first area include:

- **Global economic forums and strategic agendas:** AI has become a recurring centerpiece in global economic and political forums (e.g., Davos), where leaders debate its impact on productivity, inequality, labor markets, and global competitiveness. Its presence alongside macroeconomic and security topics reflects a shared recognition: AI is no longer a sector-specific issue, but a systemic force shaping national growth strategies and global coordination efforts.
- **AI diplomacy and international coordination:** AI has emerged as a new domain of diplomacy, with governments negotiating standards, safety frameworks, and norms for responsible use. Multilateral initiatives and bilateral agreements seek to prevent uncontrolled escalation while preserving national interests, highlighting AI's role as both a cooperative and competitive force on the global stage.
- **Direct relationships between politicians and tech leaders:** The growing strategic importance of AI has led to unprecedented direct engagement between political leaders and executives of major technology companies. These relationships influence regulation, national strategy, and crisis response, blurring traditional boundaries between public policy and private innovation.
- **Innovation-first versus regulation-first approaches:** Different regions adopt contrasting philosophies toward AI governance, with some prioritizing rapid innovation and others emphasizing precaution and regulation. These choices shape competitive positioning, market access, and ethical standards, creating divergent AI ecosystems rather than a single global market. For example, the EU has opted for binding, risk-based regulation with the EU AI Act, the United States has relied on federal executive action that can shift with each administration (such as the October 2023 Executive Order on AI, later replaced as priorities changed), China has issued dedicated regulations for generative AI services, and the UK has so far favored a pro-innovation,



principles-based approach. If you operate across several of these jurisdictions, you are effectively navigating conflicting regulatory philosophies at once.

- **Access to data and data governance:** Data remains a foundational input for AI, making access, ownership, and cross-border data flows deeply political issues. Governments increasingly regulate how data is collected, stored, and shared, balancing innovation incentives with privacy, security, and national interests. These decisions directly shape where and how AI systems can be trained and deployed.

The second area is competition between regions and states, where AI capability increasingly translates into economic and strategic power:

- **Leading geographies and the US–China dynamic:** The AI race is often framed through the lens of competition between the United States and China, each leveraging different strengths. Concretely, innovation ecosystems and private capital on one side, centralized strategy and scale on the other. This rivalry influences investment flows, trade policies, export controls, and alliance-building, affecting how AI technologies are developed, deployed, and shared worldwide. The DeepSeek case was a clear example of these competitive dynamics: In early 2025, this Chinese AI lab released open-weight models that performed competitively with leading Western systems while reportedly being trained at a fraction of the expected cost, challenging widespread assumptions about the compute and capital required to stay in the race.
- **Tech sovereignty as a national priority:** Many countries increasingly view AI as a matter of technological sovereignty, aiming to reduce dependence on foreign platforms, models, and infrastructure. This drives investments in local AI ecosystems, national cloud initiatives, and domestic talent development, reshaping procurement strategies and cross-border collaboration.
- **Multiple regions competing to join the AI race:** Beyond the US and China, regions such as Europe, the Middle East, and parts of Asia are positioning themselves as AI hubs through targeted investments, talent programs, and regulatory frameworks. This multipolar competition fragments the AI landscape while accelerating innovation through regional specialization.
- **Impact on national security and cybersecurity strategies:** In some contexts, AI is increasingly discussed alongside critical national security assets. AI enhances capabilities such as threat detection, anomaly analysis, simulation, and decision support, but it also lowers the barrier for sophisticated cyberattacks, disinformation campaigns, and automated vulnerability exploitation. This dual-use nature forces governments to rethink defensive and offensive postures simultaneously. For companies operating critical infrastructure or sensitive data, national security considerations

increasingly translate into regulatory requirements, compliance obligations, and closer alignment with state-led cybersecurity strategies.

The third area covers resources and infrastructure, the physical and financial foundations that determine who can actually build and run AI at scale:

- **Energy consumption and data center infrastructure:** AI's computational demands have elevated energy availability and data center capacity to strategic concerns. Nations with access to affordable, reliable energy and favorable infrastructure conditions gain a structural advantage, linking AI development to broader energy and climate strategies. The increasing race for bigger and more powerful model techniques is even resuscitating the appetite for traditional sources of energy such as nuclear plants and raising concerns around energy consumption and carbon footprint.
- **Semiconductor and chip manufacturing:** Advanced chips are the backbone of modern AI, making semiconductor supply chains a focal point of geopolitical tension. Export controls, domestic manufacturing incentives, and strategic stockpiling reflect the recognition that chip access directly translates into AI capability and economic power. It is impossible to discuss this topic seriously without mentioning **Taiwan Semiconductor Manufacturing Company (TSMC)**, which manufactures the vast majority of the world's most advanced chips; this concentration makes the situation around the Taiwan Strait arguably the single biggest geopolitical risk for global AI infrastructure. Also, AI infrastructure depends on critical raw materials used in chip manufacturing, batteries, and data center construction. Control over these resources and the countries that produce them adds another layer of geopolitical complexity, tying AI progress to mining, trade, and environmental policies.
- **Massive public and private investments:** Governments and corporations are committing unprecedented levels of capital to AI, viewing it as a long-term growth engine and strategic necessity. These investments influence national budgets, industrial policy, and global capital allocation, reinforcing AI's role as a cornerstone of economic strategy.
- **Doubts around scalability, real value creation, and the rise of strategic partnerships:** As generative AI adoption accelerates, growing questions remain about its long-term scalability, unit economics, and ability to deliver sustained value beyond early productivity gains. High computational costs, energy consumption, data constraints, and diminishing returns on model scaling challenge the assumption that bigger models alone will continue to drive proportional business impact. In response, many companies are shifting from pure capability races to strategic partnerships to share infrastructure, data, distribution, and risk to make AI economically viable at scale.

These alliances reflect a pragmatic recognition that real value creation increasingly depends not just on technological superiority, but on circular economy, ecosystem positioning, integration depth, and collaborative advantage.

Finally, there are ethical and societal aspects that cut across all the previous dimensions, and the language imbalance is perhaps the clearest example:

- **The language factor and why most AI systems are English-centric:** The dominance of English in AI systems is not accidental, but the result of historical, economic, and structural factors: early internet content, academic research, open-source communities, and high-quality labeled datasets have been overwhelmingly English-based. This creates an asymmetry in AI capability, quality, and accessibility across languages and regions. For governments and organizations operating in non-English contexts, this raises concerns about digital inclusion, cultural representation, and strategic dependence on foreign technologies. Addressing language imbalance is not merely a technical challenge, but a geopolitical and economic issue tied to sovereignty, market access, and the preservation of local knowledge and identity.

So, as an AI leader, what can you do with this information, besides just being attentive and aware of the news and context? You can technically incorporate these considerations as part of your continuous analysis exercise by leveraging traditional frameworks like PESTEL and using them as a checklist for your own company activities. The following checklist will enable you to do that:

#### Tip

##### PESTEL analysis checklist

**P — Political factors:** Use this lens to assess how government action, power dynamics, and geopolitics affect your AI strategy. Political risk is no longer external noise, and it directly shapes supplier choice, architecture decisions, and long-term viability.

- Are our AI initiatives exposed to geopolitical tensions (e.g., US–China relations, sanctions, export controls)?
- Do any of our AI vendors or cloud providers operate under jurisdictions that could change access or terms abruptly?
- Are we aligned with national or regional AI strategies where we operate (e.g., AI sovereignty programs, public funding)?
- Do our AI use cases intersect with sensitive areas (defense, public services, critical infrastructure)?
- Is there increasing political scrutiny or direct engagement with tech leaders in our sector that we should anticipate?

**Add your notes here:**

---

---

---

---

**E — Economic factors:** Focus on cost structures, capital flows, and economic sustainability of AI adoption. AI feels like OPEX, but its economics are deeply influenced by macro CAPEX realities.

- Are we clear on the long-term unit economics of our AI use cases (compute, inference, scaling)?
- How exposed are we to AI providers' pricing strategies driven by heavy CAPEX investments?
- Are we benefiting from AI-related investor interest without overcommitting to unrealistic ROI expectations?
- Do strategic partnerships improve our economic position versus building everything internally?
- How resilient is our AI roadmap in a downturn or funding contraction scenario?

**Add your notes here:**

---

---

---

---

**S — Social factors:** Assess cultural, workforce, trust, and societal implications. Social acceptance and internal culture will determine adoption speed more than technical readiness.

- How does AI adoption affect employee roles, skills, and trust within the organization?
- Are we prepared for internal democratization of AI beyond technical teams?
- Do language limitations (English-centric models) affect adoption, inclusion, or performance in our markets?

- How does AI usage impact our brand, reputation, and customer trust?
- Are we proactively addressing concerns around bias, fairness, and transparency?

**Add your notes here:**

---

---

---

---

**T — Technological factors:** Evaluate maturity, dependencies, scalability, and innovation risk. Technical flexibility is a strategic asset in a fast-moving, hype-driven environment.

- Are we overly dependent on a single model, provider, or architecture?
- Do we understand scalability limits (cost, latency, energy) of our generative AI use cases?
- Are we building capabilities or just consuming tools?
- How fast can we switch providers or architectures if needed?
- Are we distinguishing between experimentation and production-grade AI?
- Is our technology choice dependent on specific locations or geographies (e.g., geographic cloud regions)?

**Add your notes here:**

---

---

---

---

**E — Environmental factors:** Consider energy, infrastructure, and sustainability implications. AI strategy increasingly intersects with energy strategy and sustainability initiatives.

- Do our AI workloads significantly increase energy consumption or carbon footprint?
- Are our data center and cloud choices aligned with long-term energy availability?

- Could environmental regulations affect where and how we run AI systems?
- Are sustainability concerns becoming a constraint or differentiator in AI partnerships?
- Do we have visibility into the environmental cost of scaling AI initiatives?

**Add your notes here:**

---

---

---

---

**L — Legal factors:** Review regulatory exposure, compliance, and liability. Legal readiness determines how fast you can safely scale AI initiatives.

- Are we compliant with current and upcoming AI regulations (e.g., the EU AI Act, US federal and state-level requirements, or China's generative AI regulations)?
- Do we know where liability sits when AI systems fail or hallucinate?
- How are data access, data residency, and cross-border transfers regulated in our markets?
- Are our contracts with AI providers clear on IP, data usage, and responsibility?
- Do cybersecurity and national security requirements impose additional obligations?
- Do we operate across jurisdictions with conflicting AI regulatory approaches (e.g., regulation-first in the EU versus more innovation-first stances in the US and the UK)?

**Add your notes here:**

---

---

---

---

Same as any other strategic analysis, use this PESTEL framework to identify key risks that may affect your AI initiative. This goes beyond what you and your company do and may impact any great implementation if things go wrong with any of the PESTEL dimensions. Keep an eye out and make sure you spot key concerns upfront, or as soon as they appear in this highly changing world.

## Key takeaways

Before we close the chapter, here are the essential ideas you should retain:

- AI history shows recurring hype cycles. Treat them as a leadership manual for the current wave, not as a curiosity.
- Enterprise AI succeeds or fails for business reasons (maintenance costs, unclear ownership, unmet expectations) more often than for purely technical ones.
- Generative AI is different mainly because of accessibility, not just capability: language as the interface, minutes to adopt, OPEX-style consumption, and radically cheaper experimentation.
- Geopolitics now directly shapes AI strategy, from chip supply and energy to conflicting regulatory philosophies across jurisdictions.
- Use the PESTEL checklist as a living instrument that you revisit regularly, not as a one-time exercise.

## Summary

In this chapter, you traveled from the earliest dreams of intelligent machines to the generative AI era. We covered the main historical phases of AI and the business lessons hidden in each of them, the milestones that actually changed industry adoption, and the reasons why generative AI transforms your leadership game: lower adoption barriers, radically cheaper experimentation, a new talent equation, and new categories of risk. We also compared this wave with the printing press, extracted practical lessons from past hype cycles, and reviewed the geopolitical forces (chips, energy, regulation, and sovereignty) that now shape every AI strategy.

In practice, you can use this material in three ways: revisit the self-reflection boxes to shape your own leadership style, apply the hype-cycle lessons whenever a new technology wave reaches your desk, and run the PESTEL checklist periodically to spot external risks before they hit your roadmap. As an AI leader, you will need to divide your attention between these external signals and the internal progress of your organization.

In the next chapter, we'll start building the technical side of your AI literacy, equipping you with the machine learning knowledge that will help you connect your initiatives from the strategic and tactical levels to your technical implementations and operations.

# 3

## Exploring Key AI Terminology and Concepts

Depending on your current level of AI literacy, your path for upskilling may look a bit different. While you build your foundational knowledge for your generative AI activities, you will also need to understand a series of key terms and non-GenAI technologies that will be relevant in the context of your company's AI adoption pattern.

Regardless of your preliminary level of knowledge, this chapter will help you develop your technical vision, including a deep dive into machine learning and deep learning, before we move on to the technical foundations of language models, generative AI, and agentic solutions in general in the next chapter.

The key topics that we will cover in this chapter are:

- The eternal debate on AI's levels of intelligence
- Main AI capabilities
- AI learning
- The AI lifecycle
- Technical operations
- Deep dive into ML models
- Deep dive into DL models
- Illustrative use cases



**Note****Download the PDF version of this book**

Your purchase includes a DRM-free PDF copy of this book, the code bundle, and a range of exclusive benefits. To unlock everything, follow the *Free benefits with your book* section in the *Preface*.

## Understanding the levels of AI

General discussions around AI often revolve around levels of intelligence, rather than individual products or models. This framing usually helps researchers, policymakers, and industry leaders reason about what current systems can do, what future systems might look like, and what kinds of social, economic, and ethical consequences may follow.

While the terminology is not fully standardized, a three-level classification has become widely accepted as a useful mental model: **narrow AI (NAI)**, **artificial general intelligence (AGI)**, and **artificial superintelligence (ASI)**. But these categories are not simply technical milestones with specific requirements and benchmarks. Instead, they reflect deeper philosophical questions about intelligence, learning, and autonomy, and they shape how investments, regulations, and public narratives around AI evolve.

The reality is that these levels don't necessarily change your current AI strategy or the specific projects you may be launching, but they help you as a leader in three main aspects:

- Understanding and following industry discussions
- Keeping in mind potential and/or future risks in your current scenarios
- Adapting your AI strategy proactively for what will be next

Once again, there is no standard model for these levels of intelligence, but let us look at some general descriptions that will support your own mental framework.

### Narrow AI (NAI)

NAI refers to systems designed and trained to perform specific tasks within well-defined boundaries. This is the form of AI that dominates the industry today and underpins most real-world applications already embedded in daily life, including machine learning, natural language processing, and deep learning applications, among others. The key characteristics are:

- **Task specialization:** NAI systems are optimized for a particular function, such as image recognition, language translation, fraud detection, recommendation engines, or route optimization. They can outperform humans in narrow domains but cannot transfer that competence elsewhere.

- **Data-bound learning:** These systems rely heavily on historical data and statistical patterns. Their "understanding" is derived from correlations in training data rather than from lived experience or grounded world models.
- **Lack of general reasoning:** While modern systems may appear to reason, their behavior is fundamentally shaped by pattern matching and probabilistic inference. They do not independently form goals, develop intent, or reason abstractly outside the scope of their training.
- **No autonomous learning:** Improvements typically require retraining, fine-tuning, or architectural changes made by humans. Contrary to what many people have been saying (and selling) for years, Narrow AI does not meaningfully redefine itself or its objectives.

This category includes traditional virtual assistants, text generation for search boxes, medical imaging tools, predictive maintenance systems, and algorithmic trading platforms. Despite their impressive performance, these systems remain constrained by design. Their intelligence is instrumental and focused on single tasks, not "general". And by general, we mean that you cannot get these applications and models to learn about specific scenarios and then generalize that knowledge to approach other situations.

## Artificial General Intelligence (AGI)

The second level represents a qualitative leap rather than a simple scaling of existing capabilities. AGI is typically defined as a system that can understand, learn, and apply knowledge across a wide range of domains, adapting to new problems in ways comparable to human cognition. This is where things get tricky, as current discussions try to make sense of generative AI and LLMs to understand if we can really talk about AGI at this point, or if we are still in the NAI era. Some of the aspects often associated with AGI include:

- **Cross-domain learning and transfer:** An AGI should be able to take knowledge learned in one context and apply it meaningfully to another without task-specific retraining.
- **Robust reasoning and abstraction:** Unlike narrow systems, AGI would form internal models of the world, reason about causality, and handle novel situations with minimal prior exposure.
- **Contextual understanding:** AGI would interpret information not just syntactically, but semantically and pragmatically, understanding intent, nuance, and implications across environments.
- **Adaptive goal-directed behavior:** Such systems could dynamically adjust strategies based on feedback, constraints, and long-term objectives rather than following rigid optimization rules.

So yes, AGI is best understood as a moving target rather than a specific technical milestone. As machines absorb more tasks previously thought to require general intelligence, the definition itself evolves. Today, AGI remains an aspirational concept, though many AI researchers argue that we are seeing early precursors in increasingly general-purpose models. Some people out there may have a different feeling, as claims for AGI repeat continuously over time, but this has been a recurrent discussion for the last few decades, and it is still not over.

## Artificial Superintelligence (ASI)

ASI refers to a hypothetical stage in which machine intelligence surpasses human intelligence across all domains, including scientific creativity, emotional understanding, strategic planning, and moral reasoning. A bit like the science fiction kind of version people used to imagine during the early times of AI, and what we often get in movies and TV shows (e.g., Skynet in Terminator, Ultron in Avengers: Age of Ultron). Still far off, but a good metaphor for what a superlative AI would be, and of the considerations we would need to keep in mind. Some of the

typical attributes of ASI are:

- **Cognitive superiority:** An ASI would not only match but vastly exceed human performance in intellectual tasks, potentially discovering new scientific principles or forms of reasoning beyond human comprehension.
- **Recursive self-improvement:** Unlike human-guided optimization, ASI might improve its own architecture, learning processes, and goals autonomously.
- **Unpredictable emergent behavior:** At this level, outcomes may no longer be fully understandable or controllable by human designers.

At present, ASI remains firmly theoretical. Nonetheless, it raises profound existential questions around control, alignment, power concentration, and the future of human agency. These concerns have made ASI a central theme in both academic debate and science fiction, and they exert significant influence over long-term AI safety research and public imagination. If you feel curious about these futuristic theories, you certainly want to explore classic literature like "I, Robot" from Isaac Asimov (1950), and the famous Three Laws of Robotics we mentioned in the historical perspective of *Chapter 2*.

## The current status

So you may be wondering, where are we now? Despite rapid progress, it is broadly accepted that the world is still operating within the era of NAI. The open question is not whether today's systems are impressive, but whether continued improvements in scale, data, and architecture are sufficient to bridge the gap to AGI.

Concretely, there is a vivid debate on LLM scaling and the need for new paradigms. This is a work-in-progress discussion. On one side of the debate are researchers and companies who

believe that scaling existing approaches based on building larger and more powerful models, with more data and compute will eventually help reach AGI-like capabilities. Major AI labs and technology providers have publicly embraced this view, pointing to emergent behaviors in large models as evidence that general intelligence may arise gradually rather than through a single breakthrough. But it may be motivated (or not) by commercial or strategic interests.

On the other side are critics who argue that current architectures lack key ingredients required for true intelligence, such as top experts and Professors Gary Marcus and Yann LeCun.

Among the critics, they have consistently emphasized that today's systems lack grounded understanding of the physical and social world, struggle with causal reasoning and long-term planning, and are too attached to distributional patterns seen in training data. More recently, Yann LeCun left Meta to create their own AI lab, and defended theories based on the idea of world models. According to LeCun, intelligent systems must be able to build internal representations of how the world works by capturing causality, physics, time, and agency, instead of just relying purely on next-token prediction (like LLMs; wait for the next chapter to fully understand this) or reactive pattern matching.

These new theories include allowing an agent to simulate possible futures before acting, persistent memory and abstraction to enable reasoning across time and contexts, and goal-driven behaviors, among others.

So yes, major AI companies (including the tech providers that will call to your door, as an AI leader in your organization) tend to frame progress optimistically, highlighting generalization, multimodality, and autonomous agents as steps toward broader intelligence. And that's okay; you just need to keep discussions realistic and make sure you don't accept extreme statements of capabilities and risks, and evaluate proposed solutions and vendors in a pragmatic, well-balanced way (you will learn more about this in *Chapter 8*).

## Main AI capabilities

Let's now focus on the present. What can you actually achieve with the full set of existing AI technologies? What are the realistic capabilities you can leverage with confidence, and that will fuel your AI investments and products in the near term? Here, we present a pragmatic selection of things you can do with AI. Note that we are talking mostly about the NAI systems we discussed in the previous section, as these are the technologies that are mature enough for reliable enterprise adoption today.

For each capability, we will analyze the same four dimensions: the typical data types involved, the required infrastructure, the level of investment risk, and the most common use cases. We focus on these four aspects because, together, they answer the questions you will care about most

as a leader: what inputs you need, what it will take to build and run the solution, how likely the investment is to pay off, and where the capability is already creating value in the market.

## Forecasting

This capability focuses on predicting future numeric values based on historical data. By analyzing past sales records (such as date, product, location, or seasonality), AI models can estimate future demand or revenue. These systems typically rely on regression techniques, which are mature and well-established methods in machine learning and statistics. Here are four key dimensions that describe this capability:

- **Data types:** Structured (e.g., table, traditional databases), time-series data such as historical sales, demand, prices, seasonality indicators, and calendar information.
- **Required infrastructure:** Low to medium; traditional data warehouses and standard machine learning platforms are usually sufficient.
- **Investment risk:** Low risk, as regression and time-series forecasting techniques are mature, well-documented, and widely adopted.
- **General use cases:** Sales forecasting, demand planning, revenue prediction, inventory management, and financial projections.

## Classification

Classification involves assigning an input to a predefined category. For example, after training a model with labeled images of cats and dogs, the AI can determine whether a new image belongs to one class or the other. When there are only two possible outcomes, it is called binary classification; when there are more than two categories, it becomes multiclass classification. Here are four key dimensions that describe this capability:

- **Data types:** Labeled (i.e., annotated by human experts) structured or unstructured data such as images, text, customer attributes, or transaction records.
- **Required infrastructure:** Medium; requires labeled datasets, model training pipelines, and sometimes GPUs for image or text processing.
- **Investment risk:** Low to medium risk, largely dependent on data quality and the cost of data labeling or the availability of the required experts (e.g., doctors in healthcare organizations).
- **General use cases:** Spam detection, image recognition, churn prediction, sentiment classification, fraud detection, and document categorization.

## Detection

Detection involves identifying unusual or meaningful patterns in continuous streams of data. In industrial settings, AI can monitor signals from machines and detect anomalies that deviate from

normal behavior. This capability is often used for early warning systems, predictive maintenance, and failure prevention. Here are four key dimensions that describe this capability:

- **Data types:** Continuous or streaming data such as sensor readings, logs, signals, or system metrics.
- **Required infrastructure:** Medium to high; often involves real-time data ingestion, monitoring systems, and scalable compute.
- **Investment risk:** Medium risk, as solutions require domain expertise and careful calibration.
- **General use cases:** Predictive maintenance, anomaly detection, fraud monitoring, cybersecurity, and quality control.

## Completion

Completion refers to the ability of AI to predict what comes next in a sequence. In search engines or e-commerce platforms, this means suggesting the next word, phrase, or product as a user types. These predictions are based on learned patterns from large amounts of prior data and are commonly powered by language models. Here are four key dimensions that describe this capability:

- **Data types:** Large volumes of sequential data such as text queries, browsing behavior, or user interaction histories.
- **Required infrastructure:** Medium to high; typically requires large models, fast inference, and sometimes specialized hardware.
- **Investment risk:** Medium risk, due to model complexity and scaling considerations.
- **General use cases:** Search autocomplete, recommendation systems, query suggestions, and conversational interfaces.

## Content creation

Content creation allows AI to generate new text, images, or other media by learning from large collections of existing content. The system does not copy directly, but instead produces new material that follows similar structures, styles, or formats. This is the core capability behind modern generative AI and large language models. Here are four key dimensions that describe this capability:

- **Data types:** Large unstructured datasets such as documents, web pages, images, code repositories, or multimedia content.
- **Required infrastructure:** High; significant compute resources, model hosting, monitoring, and governance are needed.

- **Investment risk:** Medium to high risk, due to fast-moving technology, higher costs, and required expertise.
- **General use cases:** Marketing content generation, customer support responses, report writing, design assistance, and code generation.

## Automation

Automation uses AI to execute repetitive or rule-based tasks with minimal human intervention. Combined with tools like **robotic process automation (RPA)**, AI can handle workflows such as processing emails, updating records, generating reports, or managing HR tasks. More advanced approaches, including generative and agentic AI, enable systems to make decisions and take actions dynamically. Here are four key dimensions that describe this capability:

- **Data types:** Structured business data, documents, emails, process logs, and system inputs.
- **Required infrastructure:** Medium; integration with existing enterprise systems is often the main challenge (depending on the complexity of your organization and technical setup).
- **Investment risk:** Medium risk, as success depends heavily on process design and change management.
- **General use cases:** Invoice processing, HR onboarding, customer service workflows, email handling, and report generation.

## Optimization

Optimization focuses on finding the best possible outcome under given constraints. When resources such as time, budget, space, or materials are limited, AI can evaluate many possible scenarios and select the most efficient solution. This capability often builds on operational research techniques and is widely used to reduce costs and improve operational efficiency. Here are four key dimensions that describe this capability:

- **Data types:** Structured operational data, business rules, constraints, and performance metrics.
- **Required infrastructure:** Medium; often relies on optimization solvers rather than heavy machine learning infrastructure.
- **Investment risk:** Low to medium risk, as operational research techniques are mature but require specialized skills.
- **General use cases:** Supply chain optimization, routing and scheduling, resource allocation, pricing strategies, and capacity planning.

Before we conclude this section, reflect on your company's situation. Think in terms of AI capabilities and try to understand the current level of maturity and potential applications:

Tip

AI capabilities self-reflection

For each of the capabilities, try to note down examples of past projects (if any), new ideas that you may want to add to your future list of AI use cases, and the potential considerations (e.g., risks, existing gaps to reach that kind of project).

Forecast - Add your notes here:

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Classification - Add your notes here:

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Detection - Add your notes here:



---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Completion - Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Content creation - Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Automation - Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Optimization - Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

We will go into details of potential use cases and technologies later, after we review more technical details. For now, keep in mind these capabilities and preliminary applications.

You are probably starting to think in terms of resources (e.g., infrastructure, human, data) to make them happen, and that's good, as it will give you a general idea of the potential gaps and most feasible options given your current context. And maybe some of them are already happening (e.g., a search box with completion features in your intranet, an automated forecast dashboard from your sales software provider).

Now that you know what AI can realistically do for your organization, the next natural question is how these systems actually acquire those capabilities. In the next section, we will look at how AI models learn from data, which will help you understand the effort and resources behind every capability we have just reviewed.

## AI learning

There is another factor you will need to know if you want to understand what AI is nowadays: How do AI models learn? What is the training process behind that learning? And how does this notion impact modern generative AI models and the way to adopt them in your company?

In this section, we will discuss the fundamental learning patterns and will also briefly discuss data splitting, which is a key aspect in preparing training data.

## Types of learning patterns

Here are the main learning patterns you need to keep in mind:

- **Unsupervised learning:** In this approach, models are trained on unlabeled data, meaning there are no predefined answers. The objective is to discover underlying patterns, structures, or relationships directly from the data. Common tasks include clustering and dimensionality reduction, such as grouping customers based on purchasing behavior or simplifying large datasets. A typical example is clustering similar news articles by topic or grouping images based on visual features like color or texture.
- **Supervised learning:** This type of learning relies on labeled data (e.g. manual annotations from human experts), where each input is paired with a known output. The model learns to map inputs to outputs so it can make accurate predictions on new, unseen data. Supervised learning is widely used for tasks such as classification and regression. A classic example is image classification, where a model is trained to distinguish between categories like cats and dogs, and for that to happen there is an expert who labels the animal pictures with the corresponding species. Common applications include spam detection or medical diagnosis prediction.
- **Reinforcement learning:** Reinforcement learning involves an agent that interacts with an environment and learns by trial and error. The agent takes actions and receives feedback in the form of rewards or penalties, with the goal of maximizing long-term rewards. A well-known example is game-playing AI such as AlphaGo, where the system learns winning strategies through repeated gameplay. Reinforcement learning is also used in robotics, autonomous navigation, and self-driving cars, where systems learn optimal behaviors by interacting with dynamic environments.
- **Self-supervised learning:** This approach allows models to learn from unlabeled data by automatically generating labels from the data itself. Typically, the model is trained to predict part of the input using other parts, enabling it to learn meaningful representations without human annotation. A common example is predicting the next word in a sentence, which helps models learn grammar and semantics from large text corpora. Self-supervised learning is foundational to modern language models and is also used in computer vision tasks such as reconstructing missing portions of images.

There are other learning paradigms based on hybrid learning approaches that combine some of them. The one you may want to keep in mind to understand the techniques behind modern generative AI models is the notion of **reinforcement learning from human feedback (RLHF)**. This hybrid approach combines reinforcement learning with direct human input explaining what "good" behavior looks like. Instead of relying solely on environment-based rewards, the model is guided by human preferences, ratings, or comparisons of outputs. These human judgments are

used to train a reward model, which then helps align the system's behavior with human expectations. RLHF is widely used to improve the safety, usefulness, and alignment of large language models, especially in conversational and generative applications like ChatGPT or Gemini.

## Data splitting

Especially in the case of the supervised approaches (more about ML and DL, less about generative AI), the notion of **data split** becomes very relevant. Data splitting is a core practice during the preparation of the datasets for training purposes. Concretely, it involves taking the entire dataset and splitting it into three groups (typically 80% for training, 10% for validation, and the remaining 10% for testing).

This is illustrated in the following diagram:




Figure 3.1 – Data split for AI

Each of these dataset portions have a different role, and they must remain separated:

- **Training data:** It is the largest portion and is used to teach the model by letting it learn patterns and relationships from examples. In general terms, the idea is to have a dataset diverse enough to contain all the patterns you want the models to learn. For that reason, the split between training and the other 20% has to maintain all the statistical properties in all three groups.

- **Validation data:** This 10% part is kept separate and is used during training to check how well the model is learning and to adjust settings or choices, such as model complexity. This helps ensure the model is not just memorizing the training data but is learning in a way that can generalize to new situations.
- **Test data:** It is used only at the very end, after training is complete. It provides an unbiased evaluation of how the model performs on completely unseen data, similar to how it would behave in the real world. Because the test data is never used to train or tune the model, it offers the most reliable measure of true performance. Together, these three splits help build models that are both accurate and robust, reducing the risk of overfitting and misleading results.

But keep in mind, these technical tasks are the ABCs of any data and AI team. For example, your data engineers and scientists will work in a standardized manner to prepare these data samples and use them to train models, when applicable. And they will perform other tasks like **feature engineering** (i.e., choosing the relevant variables or columns to train a model) and exploratory data analysis (e.g., understanding the statistical dimensions of the dataset, quantifying missing values, validating business hypotheses based on the data). We will come back to these data-related practices when we discuss the implementation of your AI strategy in *Part 2* of this book.

These learning patterns and data preparation practices do not happen in isolation; they are part of a broader, end-to-end process. Let's now zoom out and look at the full journey of an AI solution, from the initial idea to production and continuous operations, through the lens of the AI lifecycle.

## The AI lifecycle

Another important concept is the notion of the AI lifecycle. The lifecycle describes the end-to-end process of designing, deploying, and sustaining artificial intelligence applications. It provides a structured approach to ensuring that AI systems remain accurate, efficient, and aligned with business objectives over time, especially once they are productized.

The lifecycle typically includes several core stages, ranging from data and model preparation to monitoring and continuous improvement. While these stages can be visualized in multiple ways, from a step-by-step perspective, the following are the key phases to consider:

1. **Creation of the potential AI use case and initial exploration:** Before even starting the project, you need to identify a business problem or need that AI can solve, define the specific goals and the scope of the task, assess feasibility, and explore potential datasets and methodologies. This may include hypotheses for usage scenarios (that will also impact the cost estimation), good enough performance metrics, the preliminary list of suitable data and the corresponding exploratory data analysis to better understand the

information from a statistical and quality point of view. This stage is critical to minimize any potential risk later.

2. **Data preparation as input for AI models:** In this stage, the goal is to create a data store, or the combined dataset, that will serve as input for the AI models. This depends on the nature of the AI application, and it can include tabular data, images, text, etc. It includes collecting, cleaning, and preprocessing the pre-existing data, and the goal is to ensure quality, consistency, and relevance for model training. This includes keeping in mind the volume of information required, preparing the specific datasets used for both training and testing purposes, and defining the number of data samples you will need for specific applications (e.g., 200 pictures of each animal for an automated classification system). There are data-centric approaches that give a lot of importance to the potential variants and iterations of input datasets as a way to improve performance during the experimentation phase, rather than focusing just on model adjustments.
3. **Model selection and benchmark criteria:** The idea of this stage, that in reality goes in parallel with the previous one, is to choose appropriate algorithms depending on the specific problem defined in step 1. Also, it will include defining performance metrics to compare different models, and maybe choosing two competitive approaches to compare performance and gains. Sometimes this includes a naive approach, which is a very basic implementation with lower performance, created initially as a baseline to quantify improvements with more advanced models. It also includes model-centric approaches, complementary or alternative, that focus on tweaking model parameters as part of several iterations to reach better performance.
4. **Model and data-centric iterations:** Based on the model and data oriented approaches, the idea is to improve model performance through iterative training, feature engineering (choosing the most relevant set of characteristics), technical hyperparameter tuning, and data augmentation. A good practice to combine experimentation with successful implementation results is to define a mini-iteration roadmap with your data science team, in which you will allocate a specific amount of time and human effort to concrete iteration steps, with clearly established goals, steps, alternatives, etc. This helps reduce the uncertainty related to AI implementations, while keeping a good level of experimentation and innovation.
5. **Final model choice based on pre-established metrics:** This phase focuses on evaluating model performance using predefined quantitative metrics, and selecting the best-performing model. This requires knowing the metrics related to specific kinds of models, as they are all different, and the remaining topics we will explore in the next section, such as MLOps and LLMOps.

6. **Deployment and productization:** Integrating the selected model into production systems means ensuring scalability, efficiency, and seamless interaction with applications. This relies not only on the model itself, but also on the computing capacity, network configurations, and proper integrations that will enable good performance, low latencies, and suitable model throughput to support the entire AI application and its use.
7. **Continuous monitoring and operations:** This stage is a continuous exercise. It focuses on tracking model performance over time, detecting data drift (the situation in which the statistical properties of input data change over time, leading to a decrease in the performance of AI models), and retraining when necessary, while maintaining compliance with performance, security, and ethical standards.

In reality, the details of the lifecycle will depend on the nature of the AI solution. It is not the same as the lifecycle for generative AI applications, in which your proprietary (i.e., company-owned) data is not necessarily injected into the model as part of the training process because you rely on big pretrained models, rather than other ML, DL, and NLP applications in which you have a specific model and dataset to train.

This idea connects with the notion of technical operations, which is linked with the technical level of AI management we mentioned in *Chapter 1* and which we will be discussing next.

## Technical operations

When we talk about technical operations, we refer to the set of practices, tools, and methodologies that development and data teams use to build, deploy, and maintain software and AI systems in a reliable and repeatable way. As an AI leader, you will not perform these operations yourself, but understanding them matters: they determine how quickly your teams can move from prototype to production, how much running your AI solutions will cost over time, and how resilient those solutions will be once they face real users and real data. There are different operational domains that are relevant to technical software development and AI projects, all of them different but interconnected. Some of these domains are the following:

- **Development Operations (DevOps):** DevOps is a methodology that unifies software development and IT operations to improve efficiency throughout the application lifecycle. It focuses on automation, from continuous integration and deployment to infrastructure management and real-time monitoring. **Infrastructure as Code (IaC)** enables the programmatic definition and management of servers and networks, ensuring consistent and reproducible environments. Additionally, implementing monitoring and event logging helps detect issues early and optimize performance.

DevOps promotes collaboration between teams, breaking down silos and reducing software delivery times while enhancing system stability and scalability.

- **Machine Learning Operations (MLOps):** MLOps extends DevOps principles to the lifecycle management of machine learning models, ensuring their reliability and maintainability in production environments. It includes versioning models, automating training workflows, detecting data drift, and ensuring experiment reproducibility. MLOps also focuses on model scalability, optimizing computational resource usage, and enabling continuous updates without system downtime. By integrating monitoring tools, organizations can track model performance and adjust them dynamically. This approach ensures that AI models remain accurate, operationally efficient, and aligned with business goals over time.
- **Large Language Model Operations (LLMOps):** LLMOps specializes in managing and optimizing **large language models (LLMs)**, addressing challenges related to computational resources, latency, and response control. It incorporates techniques to reduce processing costs while maintaining output quality, such as model size optimization and efficient workload distribution. Continuous monitoring ensures model performance, detecting biases or hallucinations in generated responses. Additionally, the LLMOps methodology integrates security measures to filter unwanted content and comply with privacy and ethical regulations. This framework is essential for deploying LLMs in high-performance, enterprise-grade applications, ensuring their reliability, efficiency, and alignment with user expectations.

Before moving on, let's visualize how these operational practices fit together across the end-to-end model lifecycle. The following diagram distinguishes between the tasks owned by the AI provider or model creator and those owned by your organization as the AI adopter or implementer.

The upper part illustrates the MLOps workflow, with iterative training, parameter tweaking, evaluation, and dataset updates leading to final deployment and monitoring. The lower part illustrates the LLMOps workflow, in which a pretrained foundation model is adapted through optional fine-tuning, grounding/RAG, and prompt engineering, before evaluation, deployment, and continuous adjustments (don't worry if some of these terms are new to you, as we will explore them in the next chapter). This is shown in *Figure 3.2*:






Figure 3.2 – MLOps and LLMOps

Note that DevOps does not explicitly appear in this diagram. The reason is that DevOps is a general-purpose methodology that underpins the whole software delivery process, regardless of whether AI is involved. The diagram therefore focuses on the model-specific activities that MLOps and LLMOps add on top of those foundational DevOps practices.

During the performance of these operations, there is a fundamental aspect of automated performance monitoring that will allow your teams track and evaluate specific quantitative metrics. Keep in mind that these metrics depend on the type of model and task.

Next, we will discuss the traditional **machine learning (ML)** and **deep learning (DL)** models (keep in mind that *Chapter 4* will focus on generative AI models). You have probably heard about these acronyms in the past, but not everyone has a deep understanding of what these technologies mean or do, especially in relation to the types of AI capabilities you have explored earlier in this chapter. Let's start with a deep dive into ML.

## Deep dive into ML models

Due to the mature nature of this area, ML models are a great option for specific unitary tasks that allow, for example, the detection of specific patterns in extensive datasets, segmentation of specific cohorts of clients, or classification of binary or other situations based on previously labeled data.

While you will face plenty of situations where internal and external stakeholders may feel more excited about other kinds of technologies (e.g., deep learning, generative AI) due to their level of "technical advancement" or current market attention, ML models are the base for a lot of pragmatic and solid implementations.

From a resource point of view, ML models bring obvious advantages: data science and AI professionals usually learn these techniques during their academic journey; there are plenty of code repositories with examples that can be quickly adapted and reused; the amount of required data is lower than other techniques; and the trade-off of performance versus compute infrastructure requirements is very friendly.

In this section, we will first explore the different types of tasks and models within the ML subfamily.

### Unsupervised ML models

Let us start by discussing tasks that are based on **unsupervised** approaches. These techniques are particularly valuable because they do not require labeled data, which is often the most expensive and time-consuming asset to produce in an AI project. In practice, this means your teams can extract patterns, segments, and anomalies directly from the data you already have, lowering the entry barrier for your first ML initiatives.

#### Anomaly detection

Anomaly detection helps find atypical values or outliers from specific data patterns. ML models focus on the ability to spot this information from datasets that contain "features" (characteristics, often represented as columns) related to client, user, machine, or any other context. From a learning perspective, these techniques are mainly unsupervised, which partially lowers the entry barrier related to **subject matter expert (SME)** knowledge and the need for data annotation, but they can also be supervised if there are labeled examples of normal and anomalous instances. Obviously, there is no need for you to become an expert, but here are a couple of techniques you may want to keep in mind and revisit when you have a new project that includes anomaly detection:

- **One-class support vector machine (SVM):** One-class SVM is a type of anomaly detection model that is trained on "normal" data to learn the decision boundary of

regular observations. It is particularly effective for high-dimensional data with more than 100 features and is often used in fraud detection, system health monitoring, or identifying outliers.

- **Principal component analysis (PCA):** PCA is a dimensionality reduction technique that can also be used to detect anomalies by identifying data points that lie far from the principal components. It is useful when dealing with noise-prone data and works well in cases requiring fast detection with low computational complexity.

## Clustering

Clustering is another data-centric and unsupervised approach. This means that you can use data and apply clustering techniques to find mathematical values and distances that allow you to cluster or segment different types of groups within your data. From an AI project management perspective, it is a low-risk approach as these techniques are pretty mature and rely minimally on human knowledge (i.e., knowing how many clusters you are looking for, depending on your understanding of the business).

**K-means** is a popular unsupervised algorithm used to partition data into  $k$  clusters by minimizing the variance within each cluster. It's efficient, scalable, and useful in customer segmentation, document clustering, and image compression, where labels are unknown and pattern discovery is desired.

Both anomaly detection and clustering are very useful techniques for highly concrete scenarios and are really accessible thanks to the unsupervised-first approach, which also makes your AI projects a bit simpler.

## Supervised ML models

In contrast to what we've just discussed, a vast amount of ML projects rely on human knowledge in the form of data labeling and supervised learning. In exchange, these models can replicate that human understanding of reality, and reach great levels of performance for numerical and categorical predictions, often referred to as regression and classification tasks. Let's explore the relevant supervised approaches next.

### Regression

Think of regression as a numerical prediction. This ML method estimates the relationships between a dependent variable (the target) and one or more independent variables or features, primarily used for predicting continuous numerical values such as prices, temperature, sales, etc. It basically helps predict some potential values in time. Even if you don't need to know the details of all of them, here are a few types of regression you may hear about in your AI projects:

- **Linear regression:** Linear regression is a foundational model used for predicting continuous variables using linear relationships. It is simple, interpretable, and effective

when the relationship between inputs and output is approximately linear. Imagine some ideal and straightforward scenario to predict a student's final exam score based on hours studied and attendance rate.

- **Ordinal regression:** This type of regression is used when the target variable has a natural order but no fixed interval between values. For example, predicting customer satisfaction ratings (1 = very dissatisfied, ..., 5 = very satisfied) from delivery time and product quality, where the difference between 1 and 2, or 2 and 3, is not relevant, because they are just "categories". Summarizing, the regression models the probability of outcomes that fall into ordered categories.
- **Poisson regression:** This regression technique models count data and is appropriate when the target variable represents counts (amount) or frequency of events, such as positive numbers of purchases, clicks, or calls based on specific variables such as active users, day of the week, etc. In this case, the regression models the potential number based on historical information.
- **Fast Forest Quantile regression:** This is a tree-based ensemble method with a group of decision trees designed to predict not just point estimates but also conditional distributions (quantiles in statistics) of the target variable. For example, this could help estimate the 90th percentile (value below which 90% of the data points fall) of house prices in a city based on square footage, location, and year built. It is ideal for uncertainty modeling and risk assessment, because the decision trees handle the complexity and flexibility that this kind of quantile regression needs, without assuming a strict mathematical form for the relationship between variables and the target.
- **Fast Forest regression:** This is an implementation of random forests, which as you saw earlier, is an ensemble of decision trees that outputs the average prediction of individual trees. It is non-linear, faster to train, and robust, making it suitable for a variety of regression tasks. In this case we talk about a parallel implementation. The difference with the Fast Forest Quantile regression we just saw lies in the output: the standard Fast Forest regression predicts a single expected value, whereas the quantile variant estimates ranges or percentiles of possible outcomes.
- **Bayesian Linear regression:** This variant of linear regression incorporates prior distributions over model parameters, enabling it to model uncertainty in predictions. It's useful in probabilistic modeling and when interpretability is important, and it is based on the idea of Bayesian and conditional probabilities.
- **Boosted Decision Tree regression:** Boosted trees combine multiple weak learners (decision trees) into a strong ensemble model. They are powerful for capturing complex, nonlinear relationships and are robust to overfitting with proper regularization.

Remember that boosting builds an ensemble of trees where each new tree focuses on the mistakes of the previous ones, improving accuracy over time. Therefore, we are talking about a sequential implementation.

## Classification (binary)

Instead of a numerical prediction like the case of regression models, the goal here is to group and predict by "classes" or categories. It could be as simple as classifying emails as 'spam' or 'not spam' based on features like sender address, keywords, and message length, or as complex as predicting whether a patient has a particular type of cancer based on high-dimensional data such as gene expression levels, medical imaging data, and clinical history. Some techniques include:

- **Naïve Bayes:** A simple and fast classification algorithm that uses probability and assumes all features are independent, which in reality is rarely true, but often works well.
- **Two-class SVM:** Similar to the one-class SVMs we saw for anomaly detection, two-class SVMs find the hyperplane that maximally separates two classes in feature space. They are effective in high-dimensional spaces and are robust to overfitting.
- **Two-class Averaged Perceptron:** A simple linear classifier whose final model is the average of all the weight vectors produced during training, which makes it more stable than the plain perceptron. It's fast and lightweight, useful for large-scale or streaming data where interpretability isn't a primary concern.
- **Two-class Logistic Regression:** Following the same principle of other regression models but adapted to classification scenarios, a logistic regression estimates the probability that a given input belongs to a certain class. It's interpretable, fast, and works well for linearly separable problems and applications like spam detection or churn prediction.
- **Two-class Bayes Point Machine:** This probabilistic model approximates the average of all hypotheses that correctly classify the data. It's useful when probabilistic output and theoretical soundness (solid and well-established mathematical principles and reasoning) are important, although it's less commonly used than logistic regression.
- **Two-class Decision Forest:** An ensemble of decision trees trained via bagging. This model offers robustness, handles nonlinear data well, and is suitable for many general classification tasks. The final output is decided by combining the separated model predictions through a "vote". That vote can be based on the general number of class votes per model (hard approach), or on the highest class probability voted by the models (soft approach). Voting helps improve accuracy by leveraging the strengths of different models and reducing individual errors.

- **Two-class Boosted Decision Tree:** Based on the same pattern as the regression model, this is a powerful classifier that builds trees sequentially, focusing on correcting errors made by previous ones. It excels in accuracy and handles mixed data types well, making it often used in competitive ML tasks. From a project point of view, it is a bit more complex and requires more training effort.
- **Two-class Decision Jungle:** A variant of decision forests where the structure is a directed acyclic graph (DAG) instead of a tree. It allows for feature sharing between paths and works well for complex decision boundaries.

**Classification (multiclass)**

Similar to binary classification, the multiclass variant finds more than two classes or categories. These include:

- **Multiclass Logistic Regression:** Extends logistic regression to handle more than two classes, usually via softmax (a mathematical function that converts raw numerical scores into probabilities for each class). It is interpretable and effective for problems where classes are linearly separable.
- **Multiclass Decision Forest:** An ensemble method where each tree votes for a class, and the majority wins. It is accurate and robust for multi-category tasks like image labeling or document classification.
- **Multiclass Decision Jungle:** Similar to decision forests but with a directed acyclic graph (DAG) structure that allows node reuse, improving generalization and reducing overfitting on high-dimensional datasets.
- **One-vs-All Multiclass:** A meta-strategy that breaks down a multiclass problem into multiple binary classification tasks. It allows the reuse of any binary classifier for multiclass settings.

Before we move on to exploring DL models, let's look at the following summary table of the main ML models based on the type of tasks and considerations:

| Task type         | Model name           | When to use it                                       | Considerations                                                                |
|-------------------|----------------------|------------------------------------------------------|-------------------------------------------------------------------------------|
| Anomaly detection | <b>One-class SVM</b> | When detecting outliers in high-dimensional datasets | Requires only normal class training data; good for fraud, intrusion detection |

|                          |                                         |                                                                              |                                                                                     |
|--------------------------|-----------------------------------------|------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
|                          | <b>PCA-based anomaly detection</b>      | When dimensionality reduction is needed; detects anomalies based on variance | Fast and interpretable; assumes linear relationships                                |
| Clustering               | <b>K-means</b>                          | When grouping unlabeled data into clusters                                   | Simple, scalable; needs the number of clusters defined; assumes spherical clusters. |
| Regression               | <b>Linear Regression</b>                | For predicting continuous values with linear relationships                   | Very interpretable; weak with nonlinear patterns                                    |
|                          | <b>Bayesian Linear Regression</b>       | When you need uncertainty estimates with linear regression                   | Adds probabilistic insight; more complex to implement                               |
|                          | <b>Poisson Regression</b>               | When modeling count-based data (e.g., calls per hour).                       | Good for rare event prediction; assumes Poisson-distributed output                  |
|                          | <b>Ordinal Regression</b>               | When the target has ordered categories (e.g., satisfaction ratings)          | Captures order in outputs; less commonly supported in libraries                     |
|                          | <b>Fast Forest Regression</b>           | For robust nonlinear regression tasks                                        | Handles missing values; slower training than linear models                          |
|                          | <b>Fast Forest Quantile Regression</b>  | For predicting ranges/quantiles of outcomes                                  | Useful in risk-sensitive domains; more resource-intensive                           |
|                          | <b>Boosted Decision Tree Regression</b> | For complex nonlinear regression problems                                    | High accuracy; sensitive to overfitting; requires tuning.                           |
| Two-class classification | <b>Logistic Regression</b>              | For binary classification when interpretability matters                      | Fast and explainable; poor with nonlinear patterns                                  |

|                           |                              |                                                  |                                                                  |
|---------------------------|------------------------------|--------------------------------------------------|------------------------------------------------------------------|
|                           | <b>Two-class SVM</b>         | For high-dimensional binary classification       | Accurate; needs kernel selection; not scalable for huge datasets |
|                           | <b>Averaged Perceptron</b>   | For fast linear classification                   | Lightweight; not great for non-linear data                       |
|                           | <b>Bayes Point Machine</b>   | For probabilistic classification                 | Probabilistic outputs; less commonly used                        |
|                           | <b>Decision Forest</b>       | When needing robustness across varied data       | Ensemble method; less interpretable                              |
|                           | <b>Boosted Decision Tree</b> | For high-accuracy binary tasks                   | Often a top performer; sensitive to noise and overfitting        |
|                           | <b>Decision Jungle</b>       | For compact models with reused nodes             | Efficient; less interpretable than trees                         |
| Multiclass Classification | <b>Logistic Regression</b>   | For simple multi-class classification            | Fast, interpretable; less effective for complex data             |
|                           | <b>Decision Forest</b>       | For robust multi-class classification            | Handles large feature spaces; needs tuning                       |
|                           | <b>Decision Jungle</b>       | For multi-class tasks with DAG structure         | Compact; complexity in understanding internal logic              |
|                           | <b>One-vs-All</b>            | When using binary models for multiclass problems | Flexible; higher training overhead                               |

Table 3.1: Summary of the main ML models by task type

Remember, these ML models are mature and well-known techniques that fuel plenty of AI applications around the world. You obviously don't need to know all of them, but it is good for you as a modern AI leader to get familiar with them, at least as an input/output set of capabilities, so you can understand why your team decides to use one model for a specific use case.



## Deep dive into DL models

We often describe deep learning as a subset of machine learning, or an expansion of the traditional ML capabilities, based on the concept of **neural networks**. These neural networks are computational models inspired by the structure and function of the human brain. They consist of layers of interconnected nodes called **neurons**, where each neuron processes input data by applying a weighted sum and neuron activation functions. They are based on the idea of **backpropagation**, a recurrent training process that goes backward from the output layer to the input layer so the model continues readjusting, which is like teaching a neural network by showing it its mistakes and helping it fix them step-by-step. The basic structure of a neural network is illustrated in the following diagram:




Figure 3.3 – Representation of neural networks

They learn to perform tasks by adjusting these weights through training on data, and they form the foundation of DL, where the term "deep" refers to networks of multiple hidden or intermediate layers that can automatically learn complex features and representations from raw data inputs.

In the following sections, we will first look at some key considerations that must be kept in mind while working with DL models. We will then look at the most commonly used types of DL models.

## Key considerations

Compared to the previous set of machine learning models, DL can achieve higher performance depending on the task and conditions. But there are several key considerations:

- **They are data-intensive.** Deep learning models typically have millions (or even billions) of parameters, which makes them incredibly powerful but also data-hungry. To train these models effectively and avoid overfitting, they require large volumes of labeled examples, especially in supervised settings. Without enough data, the model may memorize noise instead of learning meaningful patterns, resulting in poor generalization to new inputs. This is why industries adopting DL often invest heavily in data collection, labeling, and augmentation.
- **They make things simpler.** One of deep learning's greatest advantages over traditional machine learning is its ability to **automatically learn and extract features** from raw data. Instead of manually engineering features (like detecting edges in images or counting keyword frequencies in text), DL models learn hierarchical representations, starting with low-level patterns and combining them into increasingly abstract concepts. This automation significantly reduces the need for domain-specific feature engineering, speeding up development and often improving performance.
- **Explainability is not their forte.** Despite their power, deep learning models are often considered "black boxes" because their decision-making process is buried in many layers of non-linear transformations. Unlike simpler models (e.g., decision trees) that offer clear, rule-based reasoning, it's difficult to pinpoint exactly why a deep model made a specific prediction. This lower explainability can be a drawback in regulated industries like healthcare or finance, where understanding the reasoning behind decisions is critical, and has driven research into explainable AI (XAI) methods.
- **They have high computational cost.** Training deep models often requires powerful hardware like a **Graphics Processing Unit (GPU)** or **Tensor Processing Unit (TPU)**, which can be expensive to acquire and maintain. Complex architectures and large datasets can lead to long training times, and all of this together will certainly impact your AI project approach and roadmap.

The key takeaway is that DL is not automatically the better choice. Before committing to a DL model, your team should evaluate the trade-offs: superior performance on complex data comes at the price of larger data requirements, higher computational cost, and lower explainability. In many business scenarios, a simpler ML model may deliver comparable value faster and cheaper, so ask your team to justify when the additional complexity is really worth it.

## Types of DL models

But what are the main types of DL models? Here is a selection, in which you will recognize some of the tasks (e.g., regression, classification) from the family of ML models:

- **Neural network regression:** Just as with ML, you can perform regression tasks with deep learning. Neural networks are highly flexible models that can approximate nonlinear relationships between inputs and continuous outputs. They are powerful for large datasets with complex patterns and automated feature extraction, useful for advanced scenarios such as financial forecasting and sensor data analysis.
- **Two-class neural network:** A feedforward neural network (also known as FNN, the simplest kind of neural network with multiple neurons, without any loops or cycles between nodes) optimized for binary classification tasks. It can learn and separate complex patterns in the data, making it well-suited for tasks like sentiment analysis and fraud detection.
- **Two-class locally deep SVM:** Combines machine learning SVMs with detailed, learned representations extracted from specific small regions or parts of the input dataset (also known as local deep features, obtained via DL models) to improve classification in non-linear, complex datasets. It's beneficial in situations where your team looks for accuracy on complex or high-dimensional data.
- **Multiclass neural network:** A multi-layer model to obtain probabilities over multiple classes. It excels in problems with high complexity and data variety, such as understanding visual handwriting or performing speech recognition.
- **Convolutional neural networks (CNNs):** These are specialized neural networks designed for processing grid-like data such as images (pixels) or videos (frames). They use convolutional layers (convolutions are one kind of complex mathematical operation; for more details, refer to <https://mathworld.wolfram.com/Convolution.html>) with filters that automatically detect local patterns like edges, textures, and shapes, building hierarchical feature representations. CNNs excel at recognizing spatial structures and are highly efficient, go-to models for **computer vision (CV)** tasks.
- **Recurrent neural networks (RNNs):** RNNs are designed to handle sequential data by maintaining a hidden state or internal memory that captures information from previous inputs, allowing the network to model temporal dependencies. They are widely used for tasks like speech recognition, language modeling, and time series forecasting. However, traditional RNNs struggle with learning long-term dependencies due to issues like

vanishing gradients (they become too small) during the backpropagation process, limiting their effectiveness on longer input sequences.

- **Deep reinforcement learning:** Deep RL combines deep neural networks with reinforcement learning principles to enable agents to learn complex policies for sequential decision-making tasks. These agents are not to be confused with the "AI agents" that you hear about everywhere lately, which we'll discuss in the next chapter. These are RL agents that follow some strategy to achieve a specific reward and iterate to optimize a final result. From an AI project perspective, this is not the kind of technology you will see in your day-to-day, but it's important to know what it is about because Deep RL has powered breakthroughs in areas like game playing (e.g., AlphaGo: <https://deepmind.google/research/projects/alphago/>), robotics, and autonomous systems.

Now, let's look at the following summary table of the main DL models based on the type of tasks and considerations:

| Task type                 | Model name                                   | When to use it                                           | Considerations                                                                           |
|---------------------------|----------------------------------------------|----------------------------------------------------------|------------------------------------------------------------------------------------------|
| Regression                | <b>Neural Network Regression</b>             | For large-scale data with nonlinear dependencies         | High capacity; requires careful tuning and more compute                                  |
| Two-class classification  | <b>Neural Network (2-class)</b>              | For binary tasks with complex patterns                   | Handles nonlinear features; requires labeled data and compute                            |
|                           | <b>Locally Deep SVM (hybrid ML &amp; DL)</b> | For non-linear binary classification                     | Combines local and deep features; complex to tune                                        |
| Multiclass classification | <b>Neural Network (multi-class)</b>          | For large, nonlinear multiclass tasks                    | High performance; requires GPU/TPU resources                                             |
| Image processing          | <b>Convolutional Neural Networks (CNNs)</b>  | For images, videos, and spatial data pattern recognition | Efficient at learning local features; needs lots of labeled data and compute             |
| Sequential data           | <b>Recurrent Neural Networks (RNNs)</b>      | For sequence data like text, speech, time series         | Captures temporal dependencies; can be slow to train and suffer from vanishing gradients |

|                        |                |                                             |                                                                               |
|------------------------|----------------|---------------------------------------------|-------------------------------------------------------------------------------|
| Reinforcement Learning | <b>Deep RL</b> | For decision-making in dynamic environments | Complex training with environment interaction; requires careful reward design |
|------------------------|----------------|---------------------------------------------|-------------------------------------------------------------------------------|

Table 3.2: Summary of the main DL models by task type

Depending on the approach, keep in mind that you will need to balance complexity and time to implementation. Your team needs to understand the potential tradeoffs and required training iterations, as they will impact the number of days and required compute resources allocated for the experimentation. And that, from the AI leadership perspective, means higher investment and potential cost of opportunity.

As part of your own toolkit, keep with you the model cheat sheet from Microsoft, publicly available in the Azure Machine Learning documentation (<https://learn.microsoft.com/en-us/azure/machine-learning/?view=azureml-api-2>). It illustrates the decision-making process to select ML or DL models depending on the available dataset and expected training time.

With this overview of ML and DL models in your toolkit, you are ready to see how organizations actually put them to work. In the next section, we will review a broad set of illustrative use cases across industries, so you can start identifying opportunities in your own context.

## Illustrative use cases

There are already so many implemented ML and DL cases across industries. It may be surprising for you (or not), but certain verticals have been applying these techniques during the last three decades. From an AI leadership perspective, this feels like a safe bet: "If others were able to successfully implement these use cases, it should be ok for us. Right?"

And of course, the answer depends on so many factors (including your own company context and access to talent), but it is interesting to see the impact of time, understood as the maturity period between the moment these technologies were released and the current moment of adoption. This is interesting for you as an AI leader, regardless of the kind of technologies you want your company to adopt.

Here are a few examples that will give you an idea of the types of applications across industries:

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Finance and Banking</b></p> <ul style="list-style-type: none"> <li>• Credit scoring (logistic regression, GBM)</li> <li>• Loan default prediction</li> <li>• Transaction fraud detection (isolation forest)</li> <li>• Anti-money laundering pattern detection</li> <li>• Customer churn prediction</li> <li>• Risk-based customer segmentation (k-means)</li> <li>• Market regime detection (clustering)</li> <li>• Volatility forecasting</li> <li>• Portfolio optimization (mean-variance + ML forecasts)</li> <li>• Payment anomaly detection</li> </ul> | <p><b>Retail and E-commerce</b></p> <ul style="list-style-type: none"> <li>• Customer segmentation for personalization</li> <li>• Demand forecasting for products (time series)</li> <li>• Price elasticity modeling (regression)</li> <li>• Promotion response prediction</li> <li>• Market basket analysis (association rules)</li> <li>• Inventory stockout prediction</li> <li>• Store-level sales forecasting</li> <li>• Product return prediction</li> <li>• Customer lifetime value prediction</li> <li>• Dynamic pricing using bandits</li> </ul> |
| <p><b>Healthcare</b></p> <ul style="list-style-type: none"> <li>• Disease risk prediction from tabular data</li> <li>• Length-of-stay prediction</li> <li>• Population health segmentation</li> <li>• Radiography analysis</li> <li>• Hospital no-show prediction</li> <li>• Insurance claim cost prediction</li> <li>• Anomaly detection in lab measurements</li> <li>• Treatment outcome prediction (regression)</li> <li>• Medical billing fraud detection</li> <li>• Resource utilization forecasting</li> </ul>                                               | <p><b>Manufacturing and Industrial</b></p> <ul style="list-style-type: none"> <li>• Predictive maintenance from sensor data</li> <li>• Equipment failure classification</li> <li>• Quality defect prediction (tabular)</li> <li>• Process parameter optimization</li> <li>• Yield prediction</li> <li>• Anomaly detection in production metrics</li> <li>• Supplier risk scoring</li> <li>• Root cause analysis via feature importance</li> <li>• Production demand forecasting</li> <li>• Energy usage optimization</li> </ul>                           |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Transportation and Logistics</b></p> <ul style="list-style-type: none"> <li>• Delivery time prediction</li> <li>• Route optimization using heuristics + ML</li> <li>• Fleet maintenance prediction</li> <li>• Vehicle failure anomaly detection</li> <li>• Driver risk scoring and accident prediction</li> <li>• Fuel consumption prediction</li> <li>• Traffic volume forecasting</li> <li>• Shipment delay prediction</li> <li>• Warehouse throughput forecasting</li> <li>• Written text analysis</li> </ul> | <p><b>Telecom and Utilities</b></p> <ul style="list-style-type: none"> <li>• Customer churn prediction</li> <li>• Network outage prediction</li> <li>• Network anomaly detection</li> <li>• Call drop rate prediction</li> <li>• Energy load forecasting</li> <li>• Power consumption anomaly detection</li> <li>• Equipment failure prediction</li> <li>• Customer usage segmentation</li> <li>• Fraud detection in usage data</li> <li>• Capacity planning forecasting</li> </ul> |
| <p><b>Education</b></p> <ul style="list-style-type: none"> <li>• Student dropout prediction</li> <li>• Academic performance prediction</li> <li>• Course completion likelihood</li> <li>• Student segmentation by behavior</li> <li>• Attendance anomaly detection</li> <li>• Resource demand forecasting</li> <li>• Skill mastery prediction</li> <li>• Exam score prediction</li> <li>• Student engagement clustering</li> <li>• Intervention impact modeling</li> </ul>                                             | <p><b>Marketing and Sales</b></p> <ul style="list-style-type: none"> <li>• Lead scoring</li> <li>• Conversion probability prediction</li> <li>• Customer churn modeling</li> <li>• Campaign uplift modeling</li> <li>• Price sensitivity modeling</li> <li>• Sales forecasting</li> <li>• Upsell / cross-sell prediction</li> <li>• Customer lifetime value modeling</li> <li>• A/B test outcome prediction</li> <li>• Marketing mix modeling</li> </ul>                            |

| HR and Operations                                                                                                                                                                                                                                                                                                                                                                                                                | Government, Risk, and Smart Systems                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>• Employee attrition prediction</li><li>• Hiring success prediction</li><li>• Performance rating prediction</li><li>• Workforce demand forecasting</li><li>• Absenteeism prediction</li><li>• Salary benchmarking regression</li><li>• Workforce segmentation</li><li>• Anomaly detection in expense reports</li><li>• Productivity forecasting</li><li>• Retention risk scoring</li></ul> | <ul style="list-style-type: none"><li>• Tax fraud detection</li><li>• Benefits eligibility prediction</li><li>• Crime hotspot prediction (spatial ML)</li><li>• Infrastructure failure prediction</li><li>• Budget overrun prediction</li><li>• Traffic congestion forecasting</li><li>• Environmental pollution forecasting</li><li>• Disaster impact estimation</li><li>• Public service demand forecasting</li><li>• Policy impact simulation using ML</li></ul> |

Table 3.3: Illustrative ML and DL use cases across industries

As a continuation of your previous AI capabilities exercise, leverage this illustrative list of 100 use cases and try to see if any of them are relevant and applicable to your organization. Focus on the key concept, the input/output idea, and see if they can be interesting for any internal or external (client-facing) implementation.

Tip

Spotting Horizontal AI Applications

Note down five to ten use cases from the previous list and explain how they would be suitable for your company's activities and applications. Focus on the type of data and subfamily of ML or DL you think your team would need to implement to make them happen.

Use case #1:

Use case #2:



---

**Use case #3:**

---

---

---

**Use case #4:**

---

---

---

**Use case #5:**

---

---

---

**Other use cases:**

---

---

---

---

---

---

---

---

---

---

By now, you have started to draft potential ideas or even specific AI applications based on some initial information around AI capabilities and specific machine learning and deep learning technologies.

You can now start thinking in terms of AI use cases. The idea of use case includes, among others, the expected input data and generated outcome, as well as the type of technology and the potential business value.

One of the ways to illustrate that is by using canvases, similar to the well-known **Business Model Canvas (BMC)**, that show all the relevant information in a simple one-pager, with an easy-to-understand format that allows quick modifications and improvements. For that matter, you can leverage Louis Dorard's ML Canvas (<https://www.ownml.co/machine-learning-canvas>) as a visual you can use during your discussions with other executives, clients, and even your technical team. You can access the latest version of the canvas directly from that website and keep it as a shared reference during your AI ideation sessions.

Before you conclude this chapter, take your time to complete this **ML canvas** with one of the use cases and capabilities you have previously selected. Yes, you will focus on generative AI technologies in the next chapter, but for now, try to assimilate the ideation and definition process, regardless of the type of AI you leverage for your new use cases.

To make this exercise more approachable, let's first walk through a completed example for a simple and very common use case: customer churn prediction, i.e., anticipating which customers are likely to stop using a product or service. Here is how each block of the canvas could look for an illustrative subscription-based business:

- **Prediction task:** A binary classification model that predicts whether a customer will churn in the next three months (yes/no), based on historical customer data.
- **Decisions:** The retention team decides which customers receive proactive outreach, special offers, or account reviews, based on the churn probability produced by the model.
- **Value proposition:** Reducing customer churn protects recurring revenue and lowers acquisition costs, since retaining an existing customer is usually cheaper than winning a new one.
- **Data collection:** Customer and usage data are extracted periodically (e.g., weekly) from the CRM and billing systems, and past churn events are used as the labeled examples for training.
- **Data sources:** CRM records, billing and subscription history, product usage logs, and customer support tickets.
- **Features:** The variables the model learns from, such as customer tenure, number of support tickets in the last quarter, usage frequency trends, payment delays, and plan type.
- **Building models:** The team retrains the model on a regular basis (e.g., monthly), comparing a simple logistic regression baseline against a more advanced model such as a boosted decision tree.

- **Making predictions:** Churn scores are generated weekly for all active customers and pushed to the CRM dashboards used by the retention team.
- **Monitoring:** The team tracks prediction accuracy against actual churn outcomes and watches for data drift when customer behavior changes over time.
- **Impact simulation:** Before full deployment, the team estimates the expected financial impact, for example by comparing retention results between a group of customers managed with model predictions and a control group without them.

Now it is your turn. Use the following blank canvas to describe one of your own use cases, and don't worry if some blocks remain incomplete at this stage:

|                   |                    |                   |                 |              |
|-------------------|--------------------|-------------------|-----------------|--------------|
| Prediction Task   | Decisions          | Value Proposition | Data Collection | Data Sources |
| Impact Simulation | Making Predictions |                   | Building Models | Features     |
|                   | Monitoring         |                   |                 |              |

Table 3.4: Blank ML canvas template

You may find it difficult to know some of the technical details related to impact, data features, or the monitoring process, among others. For now, note your questions down and identify potential team members who could answer them and help you adjust the details of your ML Canvas.

## Summary

In this chapter, you have built the technical vocabulary and mental models that every AI leader needs. You explored the debate around AI's levels of intelligence (NAI, AGI, and ASI), reviewed the main AI capabilities you can realistically leverage today, and learned how AI models learn from data. You also walked through the stages of the AI lifecycle, discovered the operational disciplines of DevOps, MLOps, and LLMOps, and took a deep dive into the most relevant ML and DL models, including their trade-offs and illustrative use cases across industries. Most importantly, you can now apply this knowledge to your own work: identifying and framing candidate use cases with tools such as the ML canvas, asking your technical teams the right questions about data, models, and operations, and evaluating vendors and proposals with realistic, well-balanced criteria instead of hype-driven expectations.

In the next chapter, we will build on these foundations to explore the technical concepts behind generative AI, including language models and agentic solutions, which sit at the center of today's enterprise AI adoption.



# Part 2

## From AI Understanding to Business Transformation

With the foundations in place, the focus shifts from understanding AI to putting it to work. *Part 2* explores the technologies defining the new era of generative and agentic AI, translates that understanding into a practical AI strategy, and examines how AI can move beyond incremental improvements to reshape products, operations, business models, and the organization itself.

This part contains the following chapters:

- *Chapter 4, Understanding the New Era: Generative AI and Agentic AI*
- *Chapter 5, Crafting an AI Strategy: From Vision to Execution*
- *Chapter 6, AI-First Approaches: New Business Models and Products*



# 4

## Understanding the New Era: Generative AI and Agentic AI

The previous chapters included the foundations of your own role as an AI leader in your company, and the technical knowledge required to understand some of the main events and technologies, including those from the "pre-ChatGPT" era. You can revisit them, including the templates and checklist, to continue adjusting your leadership approach and building your technical foundations.

That said, this fourth chapter will be the cornerstone of this playbook. As you continue your path as a company leader in the generative AI era, you need to master all the technical, operational, financial, and contextual questions that will impact your success in the long term.

For obvious reasons, generative AI is taking the whole scene, and it deeply impacts your AI initiative. Even if we design a technology-agnostic AI strategy in *Chapter 5*, the reality is that your company's adoption journey will never be the same.

In this chapter, we will cover the following specific topics:

- What is new in generative AI?
- From traditional NLP to generative AI
- Identifying and prioritizing AI use cases



## What is new in generative AI?

Before we go deeper, let's go through a quick overview of the main differentiators of generative AI when compared to other AI technologies. All of the following factors have a technical base and a direct impact on the way your company can leverage generative AI technologies:

- **Drastically increased scope of knowledge:** Generative AI models are often (but not exclusively) focused on language tasks. Both **large language models (LLMs)** and **small language models (SLMs)** are trained on specific datasets, which provide them with pre-injected knowledge. This is even more relevant with multimodal models that handle different types of data such as text, images, videos, etc., to both consume and generate new information in multiple formats. From an AI adoption perspective, and simplifying the concept quite a lot, imagine the simplicity of having one single model or deployment for multiple knowledge areas in your organization.
- **Adapted to a wide variety of tasks:** One key differentiator is the ability of generative AI to perform multiple tasks with the same model. From a language point of view, one single LLM can summarize, translate, reformulate, expand, detect sentiment, etc. in one single action. This is again a huge contrast when compared to traditional models, as those are usually trained for specific tasks, with a very targeted data scope. This, combined with the previous point, brings a flexible advantage to leverage the same AI deployments for multiple projects, use cases, and applications.
- **Accelerated by a more accessible implementation:** This refers to multiple aspects that make LLMs and other generative AI models easier to leverage than previous technologies. First, the ability to consume them directly as pre-built APIs that expect a query or call and directly provide an answer. That capacity allows simple integrations with existing new software developments. The second aspect is related to how we, as humans, communicate with generative AI models. The term prompt engineering refers to the ability to create prompts or requests in plain natural language, lowering the entry barrier for anyone to use tools like ChatGPT. From an innovation perspective, this opens the door to faster innovation cycles (e.g., days instead of weeks or months) and even possibilities to include non-technical members by leveraging *vibe coding* or *no-code* tools for them to experiment and to create new concepts.

Besides these factors, what is the actual impact from an innovation point of view? What is the real change *in terms of application development and interfaces*? Generative AI is bringing new possibilities to create simple interfaces for advanced functionalities. Tools like Google's NotebookLM are great examples of the simplicity that one clean interface can bring to users, as they can leverage multiple complex models combined in one single click. There are also great

opportunities for adaptive user interfaces, neural interfaces with generative avatars, etc. When you think about potential AI solutions, don't limit yourself to the typical kind of applications. *Figure 4.1* shows how generative AI enables this new generation of user interfaces, in which a simple interaction layer gives users access to a combination of complex underlying models and capabilities:




Figure 4.1 – Generative AI and the new generation of user interfaces

However, as with many other technologies, generative AI also has limitations. Generative AI is good at reproducing information based on the specific scope of the training dataset, so that means that certain topics and tasks could be out of scope, and the performance of the model wouldn't be great. Also, they are very good at spotting linguistic information by finding semantic associations between words and topics, but a bit less with mathematical tasks (even if generative AI models are getting increasingly better over time). We will come back to these limitations, and to the techniques used to mitigate them, when covering the technical concepts in the From traditional NLP to generative AI section, later in this chapter.

## Types of generative AI capabilities

What are the types of capabilities you can obtain from generative AI technologies? How can you transform interfaces, applications, and your company's business in general? Some of these capabilities, such as text generation or summarization, will already be familiar to you thanks to tools like ChatGPT, but the value of the following list is to give you a complete, structured catalog

that you can map to potential business use cases, from the well-known basics to more advanced capabilities such as multimodal generation or access to tools. Here are the main areas:

- **Text generation:** This means producing original, human-quality written content such as emails, reports, proposals, marketing copy, and narratives by predicting and structuring language based on context, enabling organizations to scale communication and content creation with consistent quality and tone.
- **Conversational interaction:** This means supporting natural, multi-turn dialog that maintains context and intent across exchanges, forming the foundation of AI assistants, copilots, and customer-facing chat experiences that reduce friction and increase engagement. In general, once regular text generation models are created, they are then "instructed" (post-trained; you will learn more about this later) to follow a conversational pattern.
- **Language understanding and intent detection:** This refers to interpreting the meaning, sentiment, and intent behind user input, allowing systems to correctly route requests, classify information, and respond appropriately across support, operations, and analytics use cases. This capability is also related to traditional **natural language processing (NLP)** models, as you will see later in this chapter.
- **Summarization:** This involves compressing large volumes of text (documents, contracts, emails, meetings, research) into concise summaries or executive briefs, saving time and helping leaders focus on decisions rather than information processing.
- **Semantic search:** This means enabling discovery of information based on meaning rather than exact keywords, allowing employees to find answers across documents, knowledge bases, and data stores even when they don't know what to search for.
- **Reasoning:** This refers to performing multi-step thinking, drawing conclusions, comparing options, and evaluating trade-offs, making LLMs useful for complex analysis, diagnostics, and decision support rather than simple Q&A.
- **Content adaptation:** This entails rewriting, translating, localizing, re-toning, or re-formatting content for different audiences, markets, and channels, increasing reuse of existing assets and improving global consistency.
- **Code generation:** This involves writing software code, scripts, queries, and infrastructure definitions based on natural language prompts, significantly accelerating development, automation, and prototyping.
- **App modernization:** This means reading, explaining, refactoring, and improving existing codebases, supporting legacy modernization, technical debt reduction, and safer change management.

- **Text to query:** This refers to translating human questions into analytical actions over structured data, enabling non-technical users to explore metrics, trends, and anomalies without needing SQL or BI expertise.
- **Access to tools:** This means interacting with external tools, databases, APIs, and enterprise systems, extending AI beyond text generation into real-world action such as data retrieval, updates, and process execution.
- **Knowledge management:** This involves combining general model knowledge with proprietary enterprise data in real time, ensuring outputs are accurate, relevant, and aligned with company-specific context.
- **Personalization:** This entails tailoring responses based on user role, preferences, history, or business context, improving relevance, trust, and productivity across users and functions.
- **Multimodal understanding:** This means interpreting and reasoning across text, images, audio, and video, enabling use cases such as document understanding, visual inspection, media analysis, and richer customer interactions.
- **Multimodal generation:** This refers to producing content beyond text, including images, videos, diagrams, presentations, and even 3D scenes for video game-like scenarios.

The list of generative AI capabilities is continuously evolving, so keep exploring potential new use cases that may appear over time. This list includes tasks that are more or less mature and available via API interfaces from companies like Anthropic, Google, or OpenAI.

## What is Agentic AI?

During your leadership journey, you will find plenty of new concepts and terms that suddenly become mainstream and seem to be the most relevant thing of the week. Think about trending topics such as ChatGPT, Deepseek, **Model Context Protocol (MCP)**, OpenClaw (formerly Clawdbot/Moltbot), Manus, etc. You may know some (or all) of them, but the reality is that all of them took the attention of people around the world for weeks and months, and the general perception on social media was that everyone was trying and experimenting, and if you didn't, you were behind. That FOMO feeling combined with an unstoppable series of news is a perfect recipe for you, as an AI leader, to feel confused and overwhelmed.

One of the main topics where the technology is still being developed is **AI agents**, or **Agentic AI**. I'll be honest with you. I bring you this example because it is a great case that mixes technology advancements, aspirational concepts, commercial interests, and, in some cases, exaggerated statements.

From a technology perspective, agentic AI surges as a pattern that allows models to connect to all sorts of remote systems (website, databases, other AI applications) by leveraging specific tools, letting the "AI agents" analyze the context, scenario, and request, and then trial-and-error multiple times until finding a correct solution to specific task requests. This can include one single agent interacting with another system, or multiple agents interacting with each other to complete different functions (e.g., departments within a company, with one agent for finance tasks, another one for marketing, etc.). Figure 4.2 illustrates these agentic AI patterns, from a single agent connecting to external systems to multiple specialized agents collaborating with each other:




Figure 4.2 – Agentic AI patterns

The way to implement these patterns depends on the goals and technology used, and they often include "orchestrators" that understand the request and then redirect the task to other expert agents, or a network setup in which multiple (three or more) agents interact with each other, etc. One important differentiation you need to master is the one between two related concepts: AI agents and workflows. While a lot of vendors and industry practitioners are mixing both concepts, the core idea of an AI agent relies on a certain degree of autonomy in the technical decision-making while trying to accomplish a specific goal. If the series of steps is already predefined, then it is not an agent but a workflow. There is nothing wrong with the idea of a static workflow, but keep this in mind while you try to make sense of what the legion of technology sellers brings to your table.

And remember, this is a highly evolving, work-in-progress area that will likely become accessible for most of the companies. For now, it includes an endless series of acronyms and technologies like A2A (Agent to Agent from Google), **Model Context Protocol (MCP)** (widely adopted implementation from Anthropic), ACP (former Agent Communication Protocol from IBM), Microsoft Agent Framework (evolved from the previously existing Semantic Kernel and Autogen implementations), Langchain, Ollama, etc. and by the time you read this, there will be other relevant technology blocks.

But until things get a bit more stable and standardized, it is clear that you cannot follow all these technologies yourself. Instead, your technical teams can track and test them in a very pragmatic way, for example by time-boxing small proofs of concept, prioritizing the options with broader industry adoption, and quickly discarding tools that lose momentum, so that technology experimentation doesn't take a big part of your budgets. Especially when these technologies are continuously changing, their code is getting deprecated, and their market relevance is decreasing or increasing overnight.

This idea of pragmatic yet careful investment and experimentation applies to all the technical pieces involved in any generative AI implementation. Let's now explore the different key building blocks you need to understand from a functional and technology perspective.

## Key building blocks

This section is intended to be a mini capsule that helps you understand the building blocks of generative AI implementations, beyond just AI models. This stack is, in itself, part of what is new in generative AI: while traditional AI projects mostly revolved around one model and one dataset, generative AI implementations combine a broader set of components that your teams will need to assemble and operate. The goal is to show in an illustrative way the functional relation between these blocks, and the technology dependency from a technology provider perspective. Obviously, this visual is relevant for scenarios in which your team will develop the new AI use cases, rather than just adopt existing Software-as-a-Service (SaaS), online platforms such as ChatGPT or Copilot, that usually encapsulate this technical complexity in the backend, in a simplified and transparent way for the end users. *Figure 4.3* shows these building blocks and how they relate to each other, from the underlying physical infrastructure and data foundations to the user-facing interfaces:




Figure 4.3 – High-level building blocks for generative AI

Concretely, each of the blocks shown in the preceding figure plays a fundamental role in building new AI applications:

- Lakehouse (data):** This is the enterprise system of record where raw and curated business data lives (transactions, documents, logs, customer data). It feeds AI with governed, high-quality information and connects upstream to ingestion systems and downstream to the AI Platform, which selects and prepares this data for training, retrieval, and analytics. There are multiple providers for this type of solution, and your technical teams may talk about tools such as Databricks, Google BigQuery, Microsoft Fabric, or Snowflake, among others.
- Vector stores:** You can think about vector stores as an index built from enterprise data that enables AI to retrieve information based on meaning (semantics) rather than exact matches (keywords). This store or database can contain a representation of the information from other data sources (e.g., database, lakehouse), and the models and orchestrators can connect and query the vector store to find AI model responses in relevant company knowledge (e.g., policies, product docs, customer history). Along with the lakehouse, these two are key data foundations for your AI initiatives.
- Memory:** Same as commercial chatbots, memory is a mechanism that preserves context across interactions (session-level, user-level, or application-level). It works alongside vector stores and the AI Platform so AI systems can personalize responses, maintain continuity, and adapt behavior over time rather than operating statelessly.

- **AI platform:** This is the foundational layer that unifies data access, model access, orchestration, monitoring, and security. It sits between enterprise data and AI models, abstracting complexity so teams can build AI capabilities once and deploy them consistently across use cases and interfaces. You may hear about Microsoft Foundry, Amazon Bedrock, and Google Vertex AI as examples of AI platforms.
- **Orchestration engine:** It is a layer that coordinates multi-step AI tasks, based on frameworks and protocols like the ones previously mentioned in the What is Agentic AI? section (e.g., MCP or A2A). The orchestration engine connects the AI Platform, vector stores, memory, models, guardrails, and APIs to make sure that the right components are invoked in the right sequence.
- **AI models:** These make up the intelligence layer from providers like OpenAI, Anthropic, Meta, Mistral, or Google. They are invoked by the applications and the orchestration engine, protected by guardrails, and connected to knowledge bases by using data from the lakehouse and vector stores. This piece gets a lot of attention due to the competition between providers and multiple benchmarks that evaluate models like OpenAI's GPT series, Anthropic's Claude, or Google's Gemini.
- **Interfaces:** These make up the consumption layer with user and development interfaces (e.g. API) where humans and systems interact with AI outputs, such as chat applications, copilots, dashboards, or embedded features in existing products. Interfaces call into API Management, which routes requests into the orchestrated AI workflows.
- **API management:** This is the gateway that exposes AI capabilities as secure, scalable services. It enforces authentication, rate limits, versioning, and cost controls while connecting interfaces to orchestration and models in a controlled, observable way.
- **Guardrails:** These are the policy and control layer that enforces safety, compliance, and business rules before and after model execution. Guardrails integrate with orchestration, models, and API management to prevent data leakage, hallucinations, policy violations, and brand or regulatory risk.
- **End-to-end governance:** This means the cross-cutting oversight layer that spans data, models, workflows, and usage. It provides auditability, compliance reporting, lineage tracking, and risk management across the entire AI lifecycle, ensuring accountability at scale. This may converge with other internal initiatives such as data governance and cybersecurity.
- **Physical infrastructure:** The underlying compute, storage, networking, and accelerators (cloud, hybrid, or on-prem) that power model training, inference, and data processing. It is the foundation for performance, reliability, and cost efficiency, and it can be deployed in multiple technical environments and even geographic locations to



serve your applications, depending on the business context. Most of the advanced innovations happen with the public cloud as an enabler, with Microsoft, Amazon, Google, Oracle, and IBM as main hyperscalers and other contenders like CoreWeave or Crusoe. And most of them run hardware (e.g., GPUs) from companies like Nvidia (one of the most valuable companies and the main hardware actor in this generative AI era), AMD, or Intel.

If your company and teams are more or less new, they will follow an adoption and upskilling process. You can see the building blocks as pieces of a puzzle that will take time to be completed. The next section describes what this incremental adoption journey typically looks like in practice, so you can situate your own organization along the way.

## Typical adoption and customization patterns

There is no single way to adopt generative AI, but since the arrival of applications like ChatGPT, companies around the world have tried to make sense of the technical foundations and the building blocks required to implement their own "GPT-kind" of engines. Imagine this as an incremental type of deployment, in which your technical teams will add new pieces over time. It'll look something like the following figure:




Figure 4.4 – Illustrative sequence of steps for generative AI implementations

The following numbered steps expand on the sequence shown in *Figure 4.4*, in more detail and with slightly more descriptive names. They represent a chronological flow that most adopter organizations have followed since the arrival of ChatGPT, and that companies starting their journey today still tend to reproduce, usually at a faster pace. This is why the steps are described

in the past tense: they reflect the observed evolution of early adopters, while remaining a practical, current roadmap for your own journey:

1. **Deploying initial generative AI instances:** Most organizations began their generative AI journey by deploying private, enterprise-controlled conversational assistants similar to early chatbot experiences. These initial deployments typically ran isolated language models in secure environments to ensure data privacy, trust, and compliance with internal policies. At this stage, the systems were standalone: no live internet access, no integration with enterprise databases or operational systems, and no source attribution. The setup generally consisted of a baseline language model connected to a simple web or application interface, primarily used for experimentation and internal productivity.
2. **Customizing:** Baseline language models are trained on large but finite datasets, representing a historical snapshot of public information available up to a certain cutoff date that varies by model version. This limitation quickly became apparent as organizations sought answers grounded in their own proprietary content, such as internal documents, images, policies, and knowledge repositories. To address this gap, companies began augmenting models with enterprise data using a variety of customization approaches, including retrieval-based techniques, fine-tuning, and hybrid methods. These approaches enabled generative AI systems to deliver more accurate, current, and context-aware responses aligned with organizational knowledge. The following sub-bullets describe the three main customization approaches your teams will likely evaluate and combine:
  - **Prompt engineering and optimization:** This is the basic level of generative AI model customization. When you implement a solution, you can send a backdoor prompt, usually known as a meta-prompt or system prompt, that allows you to guide the model to behave in a specific manner, including tone, format of the interactions, scope of the conversations, etc. This is a highly iterative process that requires multiple optimizations to obtain the intended performance, but a cost-effective solution to guide the models, as it doesn't require any re-training, re-hosting, or complex integration.
  - **Retrieval Augmented Generation (RAG):** This technique consists of combining the generative AI models with an existing knowledge base. For example, imagine you have a data lake with all relevant documents in your organization. RAG will allow you to leverage that knowledge, and the generative AI model will know where to find it depending on the topic of the request. The positive aspects of RAG are that it is model agnostic, so you can implement it once and use it with multiple models, but it requires techniques like embeddings to reach a good level

of performance. However, it doesn't imply retraining existing models or creating new ones, so it can be another effective way to customize models, often combined with the prompt engineering and optimization part.

- **Fine-tuning:** Last but not least, you can also fine-tune an existing model. This means retraining it to modify its behavior for specific complex tasks. This can be a powerful way to adjust the behavior of certain models, but it is also costly and complex at the technical level. Also, it can be an excessive effort for any regular adopter company, as that fine-tuning won't be useful if you move to new model types and versions, which is what happens very often, given the accelerated pace of innovation. For regular (non-high-tech) companies, prompt engineering and RAG are the way to go.
3. **Securing models, prompts, and data flows:** As models became connected to internal data sources, security, privacy, and safety requirements expanded significantly. Organizations introduced private networking between models, search components, and data stores (including semantic and vector databases), along with stricter access controls. Additional protections were added to defend against prompt manipulation, misuse, and unsafe outputs, including input and output filtering, topic controls, and policy enforcement. Securing the broader infrastructure perimeter also became essential as generative AI moved closer to production environments.
  4. **Exploring frameworks, agents, and integrations:** Beyond data connectivity, adopters began integrating generative AI with external systems such as CRM, ERP, analytics platforms, and local execution environments. This phase marked the transition from simple chat experiences to more complex, action-oriented AI applications. Organizations experimented with agent-based architectures, where models could communicate with tools, systems, or even other models to plan and execute tasks. A growing ecosystem of open frameworks and libraries emerged to support orchestration, tool calling, and multi-agent coordination.
  5. **Monitoring quality, behavior, and performance:** With increasing usage, organizations recognized the need for systematic monitoring of generative AI systems. Early evaluation relied heavily on human review by subject-matter experts, but this gradually evolved toward automated metrics and testing frameworks. New evaluation dimensions became central to assessing model behavior in production. This operational discipline, often referred to as GenAIOps or LLMOps, remains an active and rapidly evolving area as tooling and standards continue to mature.
  6. **Deploying at scale through AI landing zones:** Scaling generative AI beyond pilots introduced new challenges related to performance, reliability, and cost control. Organizations responded by designing standardized, reusable deployment

environments (often referred to as AI landing zones) that bundle together models, data access, security controls, orchestration, monitoring, and cost management. These setups enable high availability, geographic distribution, capacity planning, and predictable operating costs, helping teams transition from experimentation to enterprise-grade production while maintaining governance and compliance.

7. **Converging with broader enterprise infrastructure:** As generative AI adoption matured, deployments expanded across hybrid and multi-cloud environments. Organizations connected on-premises data sources to cloud-based models using secure private links, federated data lakes across environments, and distributed vector stores. In some cases, models were deployed behind corporate firewalls or at the edge to support low-latency or disconnected scenarios. Generative AI has increasingly become a native part of the overall enterprise infrastructure rather than a standalone capability.
8. **Combining multiple models and multimodal capabilities:** Organizations began orchestrating multiple models, often from different providers, selecting the best model for each task. At the same time, multimodal models capable of handling text, images, audio, and video in a unified way expanded the scope of use cases. This shift enabled richer applications across areas such as document processing, design, customer experience, and analytics, while also improving cost and performance optimization.
9. **Incorporating data and AI governance continuously:** Governance emerged not as a final step, but as a continuous and iterative discipline. As adoption increased and regulatory frameworks evolved, generative AI became a concern not only for technology teams but also for legal, compliance, risk, and audit functions. Organizations began defining policies, controls, and accountability models covering data usage, model behavior, transparency, and lifecycle management. Effective governance now spans both organizational processes and technical controls to ensure responsible and compliant AI use at scale.

Obviously, the sequence and details of these adoption steps differ a lot depending on the organization, as well as their speed, because that will depend on the company's level of maturity, technical environment, and access to talent.

Same as in *Chapter 3*, let's reflect on your company's situation. Select five capabilities and draft some potential use cases, as well as the internal considerations for adoption in the following template:

Tip

**Generative AI capabilities: Self-reflection**

For each of the capabilities, try to note down examples of past projects (if any), new ideas that you may want to add to your future list of generative AI use cases, and the potential considerations (e.g., risks, existing gaps to reach that kind of project).

**Capability 1:** \_\_\_\_\_;

**Add your notes here:**

---

---

---

---

---

---

---

---

**Fit for your organizational goals:** ☐ Low / ☐ Medium / ☐ High

**Capability 2:** \_\_\_\_\_;

**Add your notes here:**

---

---

---

---

---

---

---

---

**Fit for your organizational goals:** ☐ Low / ☐ Medium / ☐ High

**Capability 3:** \_\_\_\_\_;

**Add your notes here:**

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Capability 4: \_\_\_\_\_;

Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Capability 5: \_\_\_\_\_;

Add your notes here:

---

---

---

---

---

---

---

Fit for your organizational goals: ☐ Low / ☐ Medium / ☐ High

Let's now move to the second part of this chapter, in which you will understand key concepts related not only to generative AI but also to the natural language processing area of study, to then deep dive and explore some technical details of the end-to-end generative AI lifecycle.

## From traditional NLP to generative AI

This section covers a series of capabilities and models related to **Natural Language Processing (NLP)** and generative AI. While generative AI is not only related to linguistic tasks, it also evolves from traditional NLP technologies to LLMs, fueling applications like ChatGPT, Microsoft Copilot, or Google Gemini. You are likely to leverage these technologies for most of your new AI projects, as companies continue to progressively adopt generative AI.

### Understanding key AI concepts

Some key concepts you need to know as "general culture" for your AI projects include:

- **Tokenization:** This is the process of breaking text into smaller units, called tokens, such as words or sub-words, for easier processing by models. It is similar to the way we learned our mother tongue in school, usually splitting sentences into several pieces. Unlike classic NLP preprocessing, tokenization keeps every piece of the text, including punctuation and sub-word fragments.
- **Vectors:** From a mathematical point of view, there is a way to process regular text or image information to make it more understandable at a technical level. It is something like translating text to numbers, as "machines" and models handle numbers in a more efficient way. Vectors are numerical representations of data (e.g., tokens, entire words, images) that allow machines to perform mathematical operations on them. From a mathematical point of view, these vectors can go from 2D or 3D representations for other non-AI domains to high-dimensional vector representations (hundreds or thousands of dimensions) for these AI tasks.
- **Embeddings:** This is a related concept that represents a learned mapping that converts discrete tokens into high-dimensional, dense vectors capturing semantic relationships between them. For example, a vector for the word "cat" can look like this: [0.12, -0.34, 0.56, 0.88, -0.45], where each number is a coordinate in the embedding space. Similar words like "kitten" or "feline" would have vectors with similar values, reflecting their related meanings and allowing AI models to spot these relations to find similarities and related information. This is the idea behind modern search and Q&A engines.
- **Attention mechanism:** This is a method that allows models to weigh the importance of different parts of input data (e.g., specific parts of a text, like key intent and entities) dynamically, improving context understanding in a more efficient manner than traditional language techniques. This is highly technical, but if you want to explore this

concept, your main reference is Google's *Attention is All You Need* paper (<https://arxiv.org/abs/1706.03762>), the real foundation of the modern generative AI era.

- **Training:** While you will learn more about AI model lifecycles in the AI model training phases section later in this chapter, it is important for you to understand that the idea of model training depends on the type of model. Training discriminative models (non-generative ones, for example, ML or DL models) and training generative AI models are not the same. While for the rest of the approaches you will rely on one model along with the dataset to train them on a specific type of task, generative AI models are pre-trained by default with a massive amount of information. In this case, the training process follows a different pattern based on both pre and post-training. You will get more information in the Pre-training and Post-training sections, but keep these concepts in mind.

The skills your AI teams will need are a bit different from those needed for machine learning and deep learning technologies, which are more science-focused. You will learn more about the evolution of roles and skills in chapter 7, but it's important to know that the industry is now moving toward more engineering-focused roles. Let's now look at the list of techniques that will consolidate your NLP and generative AI knowledge. Once again, you don't need to know all the details, but this will help you if you ever go into some level of technical detail with your implementation team:

- **Text representation:** This is the process of converting human language (text) into a numerical format (vectors) that a model can process. Here are a couple of techniques widely used for text representation, way before the current era of generative AI but still relevant to understand some of the main linguistic capabilities you have seen before.
  - **Term Frequency–Inverse Document Frequency (TF-IDF):** TF-IDF transforms text into numerical features by measuring how important a word is to a document in a collection. Additionally, these features can then be used with ML models like SVMs, logistic regression, or Naive Bayes for tasks such as spam filtering, sentiment analysis, and topic classification. This approach is lightweight, interpretable, and suitable for small to medium-sized datasets. It's certainly one of the classics for any NLP-enabled project.
  - **Bag-of-Words (BoW):** This is an alternative to TF-IDF, and a simpler way to represent text data by treating each document as an unordered collection ("bag") of words, ignoring grammar, word order, and context. This means that BoW offers a raw word count (frequency) in a document, which is good for quick baseline models, and very simple text classifications.



- **Topic modeling:** The process of finding abstract topics in text, based on unsupervised techniques such as:
  - **Latent Dirichlet Allocation (LDA):** LDA is a probabilistic model used for topic modeling, which automatically discovers hidden themes or topics in a large collection of documents. It's useful for exploring large text corpora and identifying abstract themes or clusters of documents.
  - **Non-negative Matrix Factorization (NMF):** NMF is an alternative topic modeling technique that decomposes a collection of documents into a set of interpretable topics and their associated keywords. It often produces easier-to-interpret topics than LDA for small and medium-sized corpora.
- **Embedding generation:** This entails the required techniques to create the numerical vectors you just learned about. There are different kinds of models, but some of the most common are Word2Vec, GloVe, and FastText.
- **Generative Adversarial Networks (GANs):** This concept consists of two neural networks (known as generator and discriminator) that compete in a zero-sum game. This means that the generator tries to create realistic synthetic data to fool the discriminator, which attempts to distinguish real from generated data. This adversarial training leads to the creation of highly realistic images, videos, and other data types, with applications in image synthesis, style transfer, and data augmentation.
- **Transformer models:** These models revolutionize data sequence processing via self-attention mechanisms that weigh the importance of all parts of the input sequence simultaneously, which enables efficient parallel processing and better modeling of long-range dependencies. Transformers are the foundation of systems like **Bidirectional Encoder Representations from Transformers (BERT)** and **Generative Pre-trained Transformers (GPT)**, the former created by Google and the latter by OpenAI, both built on the Transformer architecture introduced by Google.

There are plenty of techniques, but this set is highly representative and includes key types of models for some of the main capabilities you have learned before. Once again, if you ever reach this level of technical detail, you can revisit this playbook and the concepts, and of course, discuss with your AI teams to learn about new types of techniques and models.

## AI model training phases

While all AI models follow the same kind of training and inference model pattern, generative AI brings a new level of complexity. Even if the details of the training phases (post and pre) are mostly relevant to model providers or companies with strong AI engineering skills, they are still useful for you to understand what the technology providers do to prepare the models that your team will consume as an API endpoint for your AI implementation.

The lifecycle shown in the following figure relies on several stages from initial model conception and parametrization to the final customization and consumption. The specifics depend on the specific type of model, version, and the specifics of the technology provider. For example, if your team consumes a proprietary model via API (such as OpenAI's GPT series or Anthropic's Claude), the provider covers the pre-training and post-training stages, and your team focuses mostly on the customization and inference parts. However, if you adopt an open-weights model that your team fine-tunes and hosts internally, much more of the lifecycle work shifts to your own engineers. *Figure 4.5* shows this end-to-end lifecycle:




Figure 4.5 – End-to-end lifecycle of AI models

These concepts are critical because they are part of a fairly unknown area of knowledge: training and customization techniques that determine the overall scope of knowledge and performance of the models you will use for your generative AI projects.

### Pre-training

This is the foundational phase where the model learns from large-scale, unstructured datasets to develop a fundamental understanding of language, grammar, and contextual relationships. This phase is typically covered by the AI model provider. Depending on the model and the specific training approach, the pre-training phase may include techniques such as the following:

- **Self-Supervised Learning (SSL):** This is a type of unsupervised learning where a model generates its own labels from the data rather than relying on human annotations. It

allows language and multimodal models to learn representations autonomously from text, images, or other modalities. Unlike supervised learning, where labeled data is required, SSL leverages inherent structures in the data. The model is trained to predict missing or occluded parts of its input (e.g., predicting masked words in a sentence). This approach brings efficiency gains and reduced reliance on human labeling, resulting in lower operational costs and faster model iteration cycles. The main drawback is that, without human-validated labels, biases or quality issues present in the raw data can be silently absorbed by the model, which is why data curation and evaluation remain essential.

This SSL process relies on infrastructure and computing power via **Graphics Processing Units (GPUs)**. The compute requirements depend mostly on the amount of information, training iterations, intended model size, etc.

- **Vectorization:** It converts raw text into numerical representations (known as embeddings) that a model can process. Compared to regular one-hot encoding, which treats words as discrete, independent symbols, the goal of these embeddings, which are the vectors resulting from the vectorization process, is to encode the semantics and relationships between words efficiently. This means that words, sentences, or documents are transformed into multi-dimensional vectors, where similar words are closer in vector space. These models are used for text similarity, semantic search, and **retrieval augmented generation (RAG)** use cases. There are different embedding techniques, depending on their vectorization scope:
  - **Word embeddings:** Word2Vec, GloVe, and FastText generate static word representations. These embeddings are dense vector representations of words that encode both semantic meaning and relationships in a continuous, high-dimensional space. They capture semantic similarity, allowing models to understand relationships such as synonymy and analogy, although the representation of a word does not change with context.
  - **Sentence embeddings:** Models like SBERT and Universal Sentence Encoder provide vectorized sentence representations, which are numerical embeddings that capture the semantic meaning of entire sentences rather than individual words. Unlike traditional word embeddings, which provide static representations for words, sentence embeddings take context into account and provide a more holistic understanding of the input text.
  - **Contextual embeddings:** These embeddings dynamically change based on surrounding words. These embeddings are generated using deep neural networks, particularly transformer-like algorithms, and can handle highly changing

contexts, polysemy (words with different meanings), and higher dimensionality to include more features to represent information in vector spaces.

- **Expanding multi-modal:** Traditional LLMs process only text, but multimodal models integrate multiple types of data, such as images, video, and audio. From a pre-training perspective, the model can learn joint representations across different modalities, allowing it to first understand and then generate diverse types of content. The main advantage is to do so in a native way, instead of combining multiple models during the integration and inference tasks. However, it requires a large-scale dataset covering multiple modalities, and aligning different data types effectively is computationally expensive. Also, multimodality brings new challenges during the inference stage, such as the increased inference latency due to the need to process multiple high-dimensional inputs.
- **Data augmentation:** Data augmentation is a data-centric way to enhance the model performance by artificially expanding the training dataset through various transformations. This may include techniques such as the following:
  - **Synthetic data:** This approach involves creating artificial text datasets to enhance training, fine-tuning, and alignment. This approach helps overcome data scarcity, reduce biases, and generate domain-specific content efficiently. Same as regular datasets, it requires techniques such as data quality control and bias mitigation.
  - **Back-translation:** This involves translating a sentence into another language and then back to generate diverse paraphrases.
  - **Synonym replacement:** It maintains the original meaning while altering sentence structure to create multiple versions.
  - **Noise injection:** It introduces minor perturbations to the text to enhance robustness against typos.

These and other data generation and augmentation techniques help reduce overfitting and make models more generalizable, and they are especially useful for minority languages where limited labeled data exists.

- **Distributed training:** Training large-scale LLMs requires distributed computing across multiple GPUs to handle massive datasets efficiently. These GPUs are easily accessible via cloud AI infrastructure, but they rely on distributed training building blocks for large-scale training. The value of this approach is clear, but it also brings potential network latency challenges that are usually solved with clear network topologies and interconnect requirements. You can implement distributed training by using different

parallelization strategies, including data, model, or pipeline-level mechanisms. Let's learn more about them:

- **Data parallelism:** This means that the entire model is replicated across multiple GPUs, and each GPU processes a different mini-batch of data. During training, each replica computes gradients independently, which are then averaged and synchronized across GPUs before updating the model parameters. This approach is effective when the model fits within a single GPU's memory but requires large-scale data processing, so it is more suitable for highly efficient SLM.
- **Model parallelism:** This requires splitting different parts of the model across GPUs to distribute computation. The model itself is split across multiple GPUs, with each GPU handling the different parts. This approach is useful when the model is too large to fit into a single GPU, as it is the case for most of the LLMs. That said, this introduces communication overhead.
- **Pipeline parallelism:** In this case, training is done by splitting into sequential stages across multiple processors. There are different stages instead of parts, each assigned to a different GPU, forming a sequential pipeline. The batches are processed in stages, passing intermediate information across GPUs. This method is commonly used in transformer-based models, where training large-scale networks requires optimizing both memory and computational efficiency.
- **Mixture of Experts (MoE):** In MoE, the model is based on multiple "expert" subnetworks, where each expert is specialized in learning different aspects of the data. During the pre-training, the model dynamically selects a subset of these experts for each input (this is based on the idea of sparse activation patterns), allowing the model to leverage different experts' specialized knowledge for different types of tasks or input features. This approach is designed to make the model more efficient by activating only a subset of the available experts at any given time, which reduces computational costs and memory usage compared to using all experts simultaneously. The main trade-off is additional complexity: routing inputs to the right experts is difficult to train well, and unbalanced usage of the experts can degrade quality if not carefully managed.

Pre-training techniques are not widely known by most of the generative AI adopters, but we covered a good selection of them in this section. Keep in mind that, except if you are directly involved in training new models, this part will be mostly covered by the AI model provider. However, you will need to pay attention to the post-training phase, as your AI team may leverage some of these techniques to adjust pre-existing models during the AI project implementation.

## Post-training

Once pre-training is complete, models require additional refinement through post-training techniques to improve efficiency, alignment, and performance, as well as adaptation to specific kinds of tasks. The following are some of the key techniques you may hear about:

- **Fine-tuning:** This is the key part of the post-training stage, as it modifies pre-trained models for specific tasks by updating their weights using task-specific datasets. You need to know that model weights are the numerical parameters that a model learns during training, and they determine how the model processes inputs to make predictions. During the fine-tuning process, the goal is to take the pre-trained model and further train it on a smaller, task-specific dataset to adapt it for a particular application. This is great because fine-tuning adjusts model weights while retaining general knowledge, making it more efficient and flexible in terms of knowledge. However, it requires curated task-specific data and additional compute, and aggressive fine-tuning can erode part of the model's general knowledge (a phenomenon known as catastrophic forgetting). There are different ways to perform fine-tuning:
  - **Regular fine-tuning:** In this case, all model parameters are updated, requiring extensive computational resources but leading to high task performance.
  - **PEFT fine-tuning:** Based on **Parameter-Efficient Fine-Tuning (PEFT)** techniques, technologies like **Low-Rank Adaptation (LoRA)** provide an optimized approach that freezes most model weights, injects small trainable matrices into transformer layers, and significantly reduces memory usage.

Fine-tuning includes the notion of model instruction. When an LLM is instructed, the model is specifically trained to understand and respond to conversational prompts, where the focus is on aligning its behavior with user instructions. This means the model is designed to follow user commands or queries more accurately, ensuring that its responses match the expected tone, format, or context of the conversation. The model can handle more complex instructions, provide precise answers, and adapt its responses based on the specific guidance given by the user.

From a chat capability perspective, a *non-instructed LLM* could still engage in conversation and generate responses, but it might not always interpret specific user requests properly, which could lead to the generation of more general or less targeted answers or require the user to provide more context. However, most of the non-instructed LLMs are made available for posterior fine-tuning, and not to be directly used.

- **Reinforcement Learning from Human Feedback (RLHF):** RLHF is a method used to fine-tune machine learning models, particularly LLMs, by integrating human feedback into the training process. After the initial pre-training process, the model is deployed in

a real-world environment where it interacts with humans. Human evaluators act as subject matter experts (SMEs) for specific domains of knowledge and industries, and provide feedback on the model's outputs, rating them based on factors like accuracy, relevance, and ethical considerations. This feedback is used to create a reward model, which assigns numerical values to the outputs based on their quality. The model is then optimized using reinforcement learning (RL) techniques, for it to learn to adjust its responses and maximize these rewards over time. This process allows the model to align its behavior with human preferences, improving its performance and reducing undesirable outputs.

- **Pruning:** It's a post-training technique used to reduce the size and complexity of a pre-trained model by removing less important or redundant parts, such as weights, neurons, or entire layers. The goal of pruning is to make the model more efficient, improving memory usage, computational speed, and reducing inference time while striving to maintain performance. This process is especially useful for deploying large models in resource-constrained environments, like mobile devices or edge computing, where hardware limitations make running full-scale models challenging. You need to keep in mind that there is a trade-off between compression ratio and model accuracy, but pruning is certainly a good option to reach higher efficiencies. It has the following steps:
  1. The process of pruning starts with identifying which parts of the model are unnecessary or contribute minimally to its output. This can be done by examining the magnitude of weights, where small weights are considered less important, or using more advanced methods such as sparsity-based pruning, where certain connections between neurons are set to zero or removed entirely. In addition to pruning individual weights, entire neurons or even layers of the model can be pruned if they don't significantly affect the model's performance. For example, some parts of a transformer model, like attention heads, can be pruned without severely impacting accuracy.
  2. After pruning, the model typically undergoes fine-tuning to recover any lost performance due to the removal of weights or neurons. This new fine-tuning process allows the model to adapt to its reduced structure while retaining most of its learning and generalization abilities. In this case, post-pruning fine-tuning is essential because pruning can sometimes introduce a loss of accuracy, and the fine-tuning helps to restore the model's effectiveness by re-optimizing the remaining parameters.
  3. Pruning is often performed incrementally, meaning parts of the model are pruned step-by-step, with performance being evaluated and re-optimized (via iterative fine-tuning) after each step. This helps mitigate the risk of overly degrading the

model's capabilities. If too many important components are removed at once, the model could lose its ability to generalize, leading to a significant drop in performance, especially in complex tasks requiring deeper reasoning.

In general, LLM pruning optimizes LLMs by eliminating unnecessary components, which makes them more manageable for real-world applications and reduces the compute requirements for deployment.

- **Distillation:** This technique trains a smaller model (e.g., SLM) to replicate an LLM's behavior, allowing efficient deployment for specific topics and tasks. Distillation helps create more efficient "student" models from larger "teacher" models with limited loss of performance. The goal of distillation is to reduce the size and computational requirements of a model while maintaining a similar level of accuracy, making it easier to deploy in resource-constrained environments with limited processing power. It is a great option to create smaller language models based on task or domain-specific scenarios, deploy on edge devices, or even serve at scale. However, it poses two main challenges: the performance gets negatively affected when distillation does not work for highly specialized teacher models, and it is not always easy to find the right balance between model size compression and performance.
- **Quantization-Aware Training:** This technique can serve both training and inference purposes, but **quantization-aware training (QAT)** helps reduce both memory and compute requirements of generative AI models by representing the model's weights and activations with lower precision, preparing the model during the training phase for efficient inference. This process involves mapping model parameters, such as weights, to lower-bit representations (there are different levels based on geometric progression values like 1, 2, 4, 8, 16, or 32 bits), to make models more efficient for deployment on hardware with limited resources, like mobile devices, edge devices, and embedded systems, without significantly sacrificing performance. After the initial quantization, the model can be fine-tuned using the quantized parameters to compensate for any loss in accuracy caused by the reduced precision. Keep in mind that most of the current models follow 8 or 16-bit patterns such as INT8 (integer) or FP16 (half precision).
- **AI Red Teaming (AI Safety):** These are safety and security techniques applied in different moments of the model lifecycle, including post-training to ensure the model behaves properly, and continuous red teaming activities to challenge and evaluate the model in different test scenarios. Please note that AI red teaming practices are quickly evolving, and they can be applied earlier during the pre-training and later during the inference and customization steps.

This overall training process includes multiple steps, and they are the base for obtaining something between good and superlative generative AI models. Even if the pre-training activities



will be covered by the technology provider, as an AI PM, you may need to explore some post-training techniques with your data scientists and AI engineers to achieve a higher level of model performance, or to adapt existing models to specific tasks and contexts.

## Inference

This part refers to the process of using a generative AI model to generate outputs or predictions based on new, unseen input data. During inference, the model takes an input prompt and processes it to produce a response, leveraging the patterns and knowledge it has learned during its pre and post-training phases. This process is distinct from end-to-end training, as the model does not update its parameters during inference, but instead simply applies the learned knowledge to make predictions.

## Customization

After the initial training stages, the customization phase adapts the model for real-world applications. The following are some of the techniques for customized inference:

- **Chunking:** This technique means breaking long documents or datasets into smaller segments to improve information retrieval with smaller, more manageable segments, or "chunks," before processing inference queries. Since LLMs have a limited context window (they can only process a certain number of tokens at a time), chunking is a way to handle long documents, improving retrieval efficiency in RAG systems and enhancing performance in tasks like summarization, question answering, and document search by complementing models with specific knowledge bases.
- **Hybrid search:** This approach combines keyword-based (lexical) search and semantic, embedding-based search to improve information retrieval. Keyword-based methods rely on exact word matches, making them precise but limited when queries use synonyms or paraphrasing, while semantic search uses vector embeddings and retrieves results based on meaning rather than exact word matches, improving recall but sometimes returning less precise results. By combining both methods, hybrid search ensures more accurate and comprehensive retrieval, making it ideal for assistants, search engines, and document retrieval in RAG systems.
- **Reranking:** After retrieving an initial set of results, reranking applies a re-ranker model to reorder the results based on deep contextual relevance. Unlike traditional ranking methods that rely on precomputed scores, these re-rankers dynamically compare the query with each document, leading to more accurate rankings, especially for complex queries. While reranking significantly enhances precision in retrieval systems, it comes with higher computational costs, so it is typically used in a two-stage pipeline where hybrid search selects candidate documents first, and reranking optimizes their order.

## Optimization

During the inference phase, there are various optimizations that help improve response quality, efficiency, and speed for real-world applications. They are all complementary, and the list includes the following:

- **Semantic caching:** This is an industry term that refers to the practice of storing and reusing the results of previous generative AI model queries based on their underlying meaning or semantic content. Instead of recomputing the entire output for every request, the system caches results for queries with similar meanings or contexts, allowing for faster responses by retrieving cached outputs for similar inputs. This is particularly useful in scenarios where identical or similar prompts are frequently requested, reducing the computational load and improving efficiency. Semantic caching relies on understanding the content and context of the inputs, enabling intelligent reuse of past results without requiring full model inference for every new query. The main risk is returning a cached answer that is subtly wrong for a new query that only looks similar, so similarity thresholds and cache expiration policies need to be managed carefully.
- **Memory handling:** Efficient memory handling in LLMs combines several innovative approaches to optimize memory usage and improve model performance. These techniques dynamically allocate memory resources, manage large key-value caches, and reduce fragmentation to ensure that memory is used effectively during inference. By separating memory encoding and retrieval processes, some methods enable the model to store and access long-term contextual information, enhancing its ability to handle tasks involving long-form content. Additionally, models are equipped with the ability to explicitly store and retrieve knowledge in structured formats, improving their capacity for reasoning and handling complex, knowledge-intensive tasks. Temporal memory management further ensures that the model can adapt and update its knowledge over time, leading to better performance in tasks that require long-term understanding and context. Together, these strategies improve memory efficiency, scalability, and responsiveness across a wide range of applications.
- **Batch parallelism:** With this technique, you can process multiple requests simultaneously for scalability. It involves processing multiple input prompts simultaneously in a single batch rather than sequentially. This approach allows generative AI models to generate multiple pieces of content (e.g., text or images) at once, leveraging the parallel processing power of hardware accelerators like GPUs. By processing inputs in parallel, batch parallelism reduces inference time, improves hardware utilization, and increases throughput, making it ideal for high-demand, large-

scale applications. However, it may introduce some latency for individual requests and require more memory, especially for complex tasks or large batches.

- **Prompt optimization:** The idea of optimizing prompts implies carefully designing them to improve model responses and compressing them to speed up LLM inference. It can also focus on analyzing and improving the provided prompt based on advanced techniques that retro-feed the original prompt request to obtain better results.
- **1-bit quantization:** Same as the quantization concept you have seen in the *Post-training* section, this is a specific type of quantization that enables highly efficient inference of 1-bit quantized LLMs on CPUs instead of GPUs, significantly reducing memory usage and computational requirements. The key industry development for this technique is Microsoft BitNet (<https://github.com/microsoft/BitNet>), released in 2024 by the Microsoft Research team.
- **Top-K sampling:** It is a technique that injects randomness into the inference process of LLMs while ensuring that the model only selects tokens from the top K most probable ones. The idea is to control text generation randomness to obtain coherent responses. When generating text, an LLM predicts the next word (or token) based on its current context. Instead of always picking the token with the highest probability, top K sampling limits the pool of candidate tokens to the top K most likely options. From this restricted set, the model randomly selects a token, introducing variability into the generated text. This randomness helps the model avoid deterministic outputs while allowing diverse and creative responses.
- **Beam search optimization:** Compared to top K sampling, beam search optimization is a more systematic, deterministic approach, aimed at optimizing the overall sequence of tokens rather than individual token choices. During inference, LLMs use this technique to explore multiple potential sequences of words at each step and select the sequence that maximizes the overall probability. Instead of selecting just the most likely next word, beam search keeps track of the top N sequences at each generation step, expanding each sequence with the most probable tokens. This allows the model to explore multiple "paths" in parallel, leading to more coherent and contextually accurate outputs, especially in complex or long-form generation tasks.
- **Container-level optimization:** These techniques offer significant advantages for LLM inference by automating GPU node provisioning, optimizing resource utilization, and reducing operational costs through the use of lower-end GPUs for distributed workloads. Additionally, there are tools that provide fine-grained control over model configurations, data security, and fine-tuning parameters, allowing for tailored deployments while maintaining privacy and security. These features collectively

streamline LLM inference, making it more efficient, cost-effective, and flexible for various AI applications.

This section covered the different complex training cycles and techniques, along with the considerations during the inference phase. It's a highly evolving area with new techniques and libraries that are starting to get some degree of standardization at the industry level.

Obviously, this goes way beyond the level of knowledge that is expected from any executive, but it will position you in a great way during discussions with technical teams. For now, feel free to note down the techniques and steps that you consider critical during any generative AI implementation in your company (e.g., AI safety to protect the models, prompt optimization to improve performance and efficiency). The following self-reflection template will help you capture and prioritize them:

Tip

Technical Generative AI Self-reflection

From this vast list of technical steps, which ones do you think are more relevant to your company's success as an AI model adopter:

Step 1: \_\_\_\_\_;

Add your notes here:

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

Step 2: \_\_\_\_\_;

Add your notes here:

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

Step 3: \_\_\_\_\_;

Add your notes here:

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

**Step 4: \_\_\_\_\_;**

**Add your notes here:**

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

**Step 5: \_\_\_\_\_;**

**Add your notes here:**

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

## Identifying and prioritizing AI use cases

Now that you have covered a significant level of technical details for generative AI and other AI technologies such as machine learning, deep learning, and natural language processing, let's structure the ideation and prioritization process in your organization.

When we talk about designing and implementing a top-down approach to AI strategy and adoption in your company, you will need to go a bit further and make sure the investments are focused on the most promising opportunities.

You will rely on your second, tactical level of AI management (e.g., product and project managers) to structure the discussions, bring all relevant stakeholders, and, even more importantly, generate and filter the use case ideas to guide your team towards the most relevant options.

### Use case discovery

During the discovery exercise, keeping things simple is ideal. First, you will be discussing potential use cases with professionals with very different backgrounds, including executives from your own company or your clients, and other professionals with very diverse levels of actual AI knowledge.

Review the initial notes from the first chapter as you have already identified some internal stakeholders. If your use cases are for a client, you need to include their key decision-makers. And if you have the possibility, you can bring the technical AI team too.

If you assume your role as the key AI leader in your organization, you will need to lead and orchestrate discussions and make sure wrong ideas are discarded with justified arguments. You cannot simply decide by yourself.

On the other hand, you can only facilitate the discovery by using simple visual references that skip any additional complexity. Remember, people are often afraid of not knowing certain AI topics, and discussions can get too complex if there is no proper structure.

These templates are exactly what I use in my AI use case discovery workshops. A simple spreadsheet or slide can help you organize the discussion case by case, and adds a high level of information during the live discussion: the what (use case name), the who (involved department, and if the use case is to be used internally or for customer-facing applications), and the envisioned type of project. The following table captures these details:

| #   | Use case                   | Department       | Int/ext case | Type of project (Generative AI, ML, DL, NLP, etc.) |
|-----|----------------------------|------------------|--------------|----------------------------------------------------|
| 1   | Support Chatbot L1         | Customer Success | External     | Generative AI, LLM + RAG                           |
| 2   | Numerical sales prediction | Sales            | Internal     | Machine Learning, regression tasks                 |
| 3   | Social media score         | Marketing        | Internal     | Natural language processing, sentiment analysis    |
| 4   | Web bot                    | All              | External     | Generative AI                                      |
| 5   | Customer segmentation      | Marketing (VP)   | Internal     | Clustering ML                                      |
| ... |                            |                  |              |                                                    |

Table 4.1: Use case discovery template with illustrative examples

The output of this step is a simple yet valuable catalog of use cases that will be completed in the next step by evaluating the potential impact and available resources.

# Use case evaluation

As a continuation of the previous table, the use case evaluation exercise tries to collect high-level information about business impact and the available resources that determine technical feasibility, which will help during the prioritization exercise. Try to clarify, for each use case during the discussion, whether the project's value stems from cost reduction, revenue generation, or company differentiation, and define the business KPIs that will sustain these value and impact hypotheses. During or after that discussion, don't forget to document the potential available resources (or the lack of them), depending on the type of AI project and the required technical, organizational, and human resources. The following table extends the previous use case catalog with these evaluation details, again with illustrative examples:

| # | Use case                   | Impact                                                            | KPIs                   | Available resources (talent, budget, sponsor, etc.)                                                               |
|---|----------------------------|-------------------------------------------------------------------|------------------------|-------------------------------------------------------------------------------------------------------------------|
| 1 | Support Chatbot L1         | Cost-saving for solving issues                                    | Saved time per day     | AI engineers, cloud infra, some internal knowledge bases, sponsored by CVP customer success                       |
| 2 | Numerical sales prediction | Better planning                                                   | Forecast accuracy      | Partial data                                                                                                      |
| 3 | Social media score         | Increasing revenue by better targeting clients                    | Client score           | No sponsor                                                                                                        |
| 4 | Web bot                    | Cost savings for customer success and increased user satisfaction | Time to solve an issue | To be defined                                                                                                     |
| 5 | Customer segmentation      | Better marketing                                                  | % of clicks            | Infra for ML, data scientists with relevant clustering experience, and VP support. We have all the data required. |

|     |  |  |  |  |
|-----|--|--|--|--|
| ... |  |  |  |  |
|-----|--|--|--|--|

Table 4.2: Use case evaluation with impact, KPIs, and available resources

By the end of this exercise, you will have a clear view of the potential use cases, and you will start analyzing data requirements and other resources in detail to get some sense of and prioritize the use cases.

Prioritization models

Finally, based on the mix of information for the different use cases, you should be able to guide the discussions and prioritize the most interesting cases for your organization.

There are different ways to do this. For example, you can leverage a regular 2x2 matrix with "business value" and "technical feasibility" axes and place your use cases to choose a potential higher value and feasibility (often represented as the top-right quadrant of the matrix). Figure 4.6 shows one such possibility:




Figure 4.6 – 2x2 matrix for use case prioritization

Assuming you have collected a few interesting use cases, you can leverage the prioritization to create a long-term, inter-use case roadmap that connects all of them, as shown in Figure 4.7:






Figure 4.7 – Inter-use case AI roadmap

From an AI leadership perspective, you can leverage this long-term plan to progressively increase the AI maturity based on project experiences. If you select the potential sequence of projects properly, you may find some "economy of scale" as you will reuse underlying infrastructure, acquired skills, and prepared data from the first projects.

Alternatively, to the 2x2 matrix (or complementary, if you use it after a preliminary prioritization with that matrix), you can leverage my very own PTRC™ model, which stands for the four key stakeholder areas involved: Product, Technology, Revenue, and Compliance. With this model, I try to get the buy-in from those different parts of the organization to guarantee these projects are not isolated initiatives. It looks like the following figure:

## Product

*Dealing with the opportunity cost (adding a new AI feature vs others)*

## Technology

*Implementing and overcoming the innovation dilemma (evolving AI models and tech)*

## Revenue

*Optimizing the impact of AI on margin and SaaS pricing (making the most of new AI features)*

## Compliance

*Innovating while reducing legal risk (adapting compliance and security to specific UI interfaces and features)*




Figure 4.8 – The PTRC model for use case prioritization and stakeholder buy-in

Concretely, you can list the top 3 or 5 projects, and then initiate a series of discussions with your stakeholders from the product, technology, revenue, and compliance areas. And you can add ok/nok or quantitative scores to each use case, to then find an agreement and understanding of which use cases and related projects are more important to your organization. The following figure shows what that could look like:

|   |            | Product | Tech | Revenue | Compliance |
|---|------------|---------|------|---------|------------|
| 1 | Use case 1 | ✓       |      |         |            |
| 2 | ...        |         |      |         |            |
| 3 | Use case 3 | 7       | 10   | 10      | 10         |
| 4 | ...        |         |      |         |            |
| 5 | Use case 5 |         | ✗    |         |            |

Figure 4.9 – Example of PTRC-based scoring for a list of AI use cases

The rest of the business units of your company (e.g., marketing, finance) are, of course, important, but these four are the first level of influence for your AI initiatives, and to enable experimentation and potential new business models (as you will see in *Chapter 6*). You can proceed with the rest of the areas once the decision is made, so they are properly informed and ready to adapt their own activities to this new era of (generative) AI project innovation.

## Summary

In this chapter, we explored what makes the current era of generative AI and agentic AI genuinely different from previous waves of AI technology. We reviewed the key differentiators of generative AI, its main capability areas, and the emerging world of AI agents, including the important distinction between agents and workflows. We then covered the key building blocks of enterprise generative AI implementations, the typical adoption and customization patterns (from initial private deployments to prompt engineering, RAG, fine-tuning, security, monitoring, and governance), and the essential concepts that connect traditional NLP with modern generative AI, including the end-to-end model lifecycle across pre-training, post-training, inference, customization, and optimization. Finally, we introduced practical templates to discover, evaluate, and prioritize AI use cases, including the 2x2 prioritization matrix and the PTRC model.

All of this will allow you to hold informed conversations with your technical teams, ask the right questions to technology vendors, and make better adoption decisions. In the next chapter, we will build on these foundations to design a technology-agnostic AI strategy for your company, connecting these technical possibilities with your business priorities.



# 5

## Crafting an AI Strategy: From Vision to Execution

So far in this book, you have been building the foundations of your role as an AI leader, from understanding the technology and market landscape to applying your first management tools, such as the *AI Usage Policy* and the *AI Management Evaluation Checklist* in *Chapter 1*. This chapter is where things get interesting for you as an AI leader, as it is time to consolidate all those elements into a company-level plan. The idea of an AI strategy has been around for a while, with technology companies and consulting firms creating their own frameworks to structure strategic workstreams and manage the (hopefully) increasing level of maturity.

The problem for many leadership teams is that these generic frameworks rarely translate into concrete action within their own organizations. For this reason, in this chapter, instead of just learning about typical workstreams and stages, you will approach different strategic activities that will allow you to consolidate your strategy and vision and make sure that it evolves over time and materializes into a proper execution.

We will cover the following specific topics in this chapter:

- The pillars of an AI strategy
- Assessing your organization's AI maturity
- Sharing your strategic AI roadmap
- Turning your AI strategy into operational actions
- Approaching the next chapters optimally

## The pillars of an AI strategy

Let's start by defining what an AI strategy is and the main areas of work. The main idea of an AI strategy is a general, long-term plan at the company level that includes all the pillars necessary to establish the foundations of AI adoption from a top-down perspective. That means that it should be concrete enough to contain actual pieces of work in your company, and widely known and accepted so it can be implemented by the different hierarchy levels, teams, and employees.

If you simply type and search "AI strategy areas" or "AI strategy streams", you will find multiple versions from different organizations approaching this topic. However, the results would be relatively similar to each other, with only small differences between frameworks. Regardless of your company's size, sector, or current level of AI adoption, the typical pillars of an AI strategy are shown in *Figure 5.1*:




Figure 5.1: Pillars of an AI strategy and relation with the next book chapters

Let's learn more about them:

1. **Business strategy:** This defines how AI aligns with the company's overall goals to create sustainable value. The long-term vision and strategic objectives of AI implementation are key because they determine where the company invests its limited resources and prevent AI efforts from becoming a collection of disconnected experiments. The AI strategy is closely tied to your company's overall business strategy and should support broader strategic objectives and lay the groundwork for effectively implementing AI

technologies (remember the application of the Porter framework in Chapter 1 to align specific AI investments to the general strategic goals).

Depending on your specific area of work as a company leader, you may be focusing on concrete topics and play a more transversal role if you are part of the C-suite team. Last but not least, this long-term vision will also impact the way you evolve your company's business model and activities, as you will see in *Chapter 6*. Regardless of your role, some important actions in this pillar include:

- **Aligning AI investments with strategic goals:** Map every AI initiative to one of the company's strategic objectives, so each project can be justified in terms of business value rather than technological novelty.
  - **Defining measurable business outcomes:** Establish the KPIs that AI initiatives are expected to improve (e.g., revenue growth, cost reduction, or customer satisfaction) before launching them, so progress can be tracked and communicated over time.
  - **Securing executive sponsorship and funding:** Assign clear ownership at the leadership level and allocate realistic, multi-year budgets, signaling that AI is a strategic priority rather than a one-off experiment.
2. **Organization and culture:** This pillar refers to how well a company is culturally prepared to adopt AI, including workforce skills and acceptance. You will explore actions pertaining to these topics in *Chapter 7*, but the important actions include:
- **Providing training in AI and other digital skills:** This entails internal programs to develop AI and digital expertise among employees to develop a data-driven and AI-ready mindset. The training can be structured at several levels, going from the executive team to the rest of the organization.
  - **Enabling continuous innovation:** This can include initiatives like hackathons, innovation labs, and incentive programs to promote experimentation. Another fundamental aspect is the notion of knowledge transfer after AI projects, which means that the implementation teams can share their learnings, challenges, and conclusions with other teams within the company, as a sort of community of practice.
  - **Forming multidisciplinary teams:** One of the best ways to achieve better results in AI adoption is to combine technical, business, and industry-related talent. This means bringing together data and AI experts, and other teams and departments to foster collaboration, ideation, and holistic planning for your AI initiatives.

3. **Technology strategy:** This pillar focuses on the technology stack (including the AI infrastructure) and data management strategies required for an effective AI implementation. For this to work, you will rely on the relationships with technology providers, as you will explore in *Chapter 8*. Some of the key elements include:
  - **Cloud platforms:** AI projects require scalable infrastructure with adaptive computing resources, pre-built AI/ML services, and secure connections between storage and AI models. Cloud adoption enables cost-efficient scaling, automated resource management, and compliance with data security regulations, ensuring reliable model deployment and accessibility.
  - **Data pipelines:** AI success depends on data, so the quality and reliability of your data pipelines directly condition the results of your AI initiatives. In practice, these pipelines include data ingestion (from databases, APIs, or sensors), data transformation (cleaning, normalization, feature extraction, etc.), storage solutions, and processing models.
  - **System integration:** AI models must integrate seamlessly with business applications through APIs, microservices architectures, and scalable computing frameworks. Distributed processing enables handling large datasets, while edge computing supports real-time AI inference in low-latency environments, ensuring efficient decision-making.
4. **Internal governance:** AI governance ensures the ethical and responsible use of AI. You will explore all the details in *Chapter 9*. Depending on the company, it can cover several points such as:
  - **Compliance with external regulations:** Standards like the EU AI Act or ISO 42001 provide guidelines for responsible AI implementation, and serve as a basis to establish the pillars of internal governance, including comprehensive checklists, processes, and tangible audits.
  - **AI ethics committees:** These committees review AI projects to ensure they align with ethical standards and organizational goals. The review process can be both top-down and bottom-up, depending on the originator of the review request. For example, the executive level of your company can assign a specific review to the committee, but the product and technology teams can also elevate a request in case of doubt.
  - **Governance frameworks, privacy, and usage policies:** This means implementing data privacy and internal AI usage policies, like the one you drafted

in *Chapter 1*. The governance process also focuses on ensuring fairness, transparency, and accountability in AI decision-making through regular audits.

5. **Overall AI maturity:** This involves the practical implementation of AI solutions, seen as an experimental and incremental exercise at the company level. Obviously, it relies on the rest of the pillars, and it is the key measure of success towards actual AI execution in your company. This last pillar focuses on several key aspects, such as:
  - **Working with AI models and tools:** Developing AI applications involves defining a clear business problem, collecting and preparing high-quality data, selecting appropriate algorithms, and training the model (when required) to recognize patterns and make accurate predictions. Once trained, models are tested and validated to ensure they perform well under real-world conditions. The final step is deployment, where the AI system is integrated into business operations, followed by continuous monitoring to improve performance over time. Successful AI model development requires a balance between technical accuracy and practical usability, ensuring that AI solutions provide measurable value to the organization.
  - **Experimentation and iteration:** AI models require continuous testing and refinement to remain effective, as business needs and data patterns evolve. Experimentation involves A/B testing, adjusting hyperparameters, analyzing errors, and incorporating user feedback to improve performance. Iteration ensures AI systems adapt to real-world complexities, allowing businesses to fine-tune their models for higher accuracy and reliability. This continuous process helps AI solutions remain relevant, scalable, and capable of delivering optimal results, reducing the risks associated with static, outdated models.
  - **Evaluating feasibility and impact:** Before investing in AI, organizations must assess technical feasibility, implementation costs, and potential return on investment (ROI). This includes evaluating whether sufficient data and infrastructure are available, estimating development timelines, and analyzing cost-benefit trade-offs. Organizations that skip this evaluation risk investing in stalled pilots, accumulating sunk costs, and losing internal credibility for future AI initiatives. AI solutions should be scalable, ensuring they can handle growing data volumes and business demands. Measuring business impact through KPIs, such as revenue increases, operational efficiency improvements, or enhanced customer satisfaction, helps justify AI investments and ensures long-term sustainability.



The evolution of your level of AI maturity is not static, like the scope and approach of your overall AI strategy. You will see in the next section, *Assessing your organization's AI maturity*, how you can estimate your current level of maturity, analyze the critical next steps, and evolve positively over time.

Remember that all these areas are interconnected and collectively contribute to a company's AI maturity. Successfully integrating AI requires a holistic strategy that goes beyond technology, addressing alignment with business goals, organizational readiness, governance, and culture.

Also, keep in mind that no company implements the entire AI strategy in one sitting. It takes iterative work, and incremental scope to continue defining and implementing new initiatives, policies, and processes. Let's now explore the ways you can structure your maturity stages to track the evolution of your AI strategy over time.

## Assessing your organization's AI maturity

Think about the classic crawl/walk/run model, a well-known progression in which organizations first learn the basics (crawl), then establish repeatable practices (walk), and finally accelerate at scale (run). What you want here is to track how things evolve in terms of AI strategy implementation and acquired internal maturity. This assessment matters because organizations that do not measure their AI maturity tend to overestimate their readiness, launch initiatives that their data, platforms, or teams cannot yet support, and struggle to justify continued investment when results fall short. Keeping in mind that the level of organizational AI maturity, as a pillar, is a consequence and combination of the other business, technology, human, and governance factors, you try to extract or make sense of the level of AI readiness and ongoing developments.

This means that you will need to track how things are going, see if there is any critical blocker from any of the pillars, spot the main areas of progress, etc. And based on that, define different stages of maturity and the main milestones you want to see before officially confirming that your company is moving to the next level.

To create your own set of stages and milestones, remember the following:

1. **Exploring:** This is the initial stage of AI maturity, where organizations begin experimenting with AI in a limited and exploratory manner. AI initiatives are typically confined to selected use cases or functions, with the primary goal of testing capabilities and building basic familiarity. This stage is a natural entry point for most adopters, where organizations move from curiosity to early experimentation.
2. **Planning:** At this stage, organizations start defining an enterprise-level AI strategy alongside an initial portfolio of **proofs-of-concept (PoCs)**. The focus shifts from isolated experiments to intentional planning, evaluating where and how AI can be deployed to support business objectives. Companies begin assessing required data,

technology, and skills, leveraging early learnings to strengthen their overall technology stack and readiness for broader adoption.

3. **Implementing:** Early adopter organizations build on their early experience by evaluating PoC results and transitioning successful initiatives into production. AI solutions start delivering tangible value, creating internal momentum and increasing stakeholder confidence. During this phase, a strong business strategy remains essential because it keeps production efforts anchored to measurable business outcomes, while technology and data strategies begin to mature and consolidate to support scalable deployment.
4. **Scaling:** In this phase, AI solutions expand beyond individual teams or functions to become embedded across the organization. Successful implementations are replicated, and value realization becomes evident across multiple business areas. Organizational alignment, change management, and culture become critical enablers, ensuring teams are equipped to adopt and scale AI-driven ways of working.
5. **Realizing:** This is the most advanced level of AI maturity. Organizations operate AI at scale through standardized, repeatable, and measurable processes. AI strategy, governance, and operating models are fully integrated into the business, enabling continuous optimization and innovation. Governance evolves alongside experience, ensuring responsible, sustainable, and value-driven AI adoption aligned with organizational goals.

These stages focus on the evolving level of the company, which builds on the growing experiences, success stories, gained efficiencies, and clear leadership support. Remember that depending on your company's context, the level of progress in each of the different pillars may be different, so your expectations will need to be realistic (e.g., you cannot move from stage 1 to 3 if your technology strategy is not ready; you cannot reach the final "realizing" stage if your business and AI strategies are not aligned).

Before we move to the next section, let's spend some time analyzing your company's case. You can even recover the notes from the *AI Management Evaluation Checklist* in *Chapter 1*, as most of those answers will be relevant for the current exercise.

To self-assess each pillar consistently, use the following reference for the three levels: select Low if the pillar is still at the level of ideas or isolated conversations, with no formal initiatives in place; select Medium if there are defined plans and first initiatives underway, but practices are not yet consistent across the organization; and select High if the pillar is supported by established, repeatable practices with clear ownership and measurable results.

Tip

**AI Maturity and Gap Assessment**

In this case, you will first evaluate the current progress for all five pillars of the AI strategy, and then reflect on the estimated level of maturity.

**1. Business Strategy**

Estimated level of progress on this pillar:

☐ Low / ☐ Medium / ☐ High

Relevant accomplishments:

Existing gaps:

Next steps:

**2. Organization and Culture**

Estimated level of progress on this pillar:

☐ Low / ☐ Medium / ☐ High

Relevant accomplishments:

Existing gaps:

Next steps:

3. Technology Strategy

Estimated level of progress on this pillar:

☐ Low / ☐ Medium / ☐ High

Relevant accomplishments:

Existing gaps:

Next steps:

4. Internal Governance

Estimated level of progress on this pillar:

☐ Low / ☐ Medium / ☐ High

Relevant accomplishments:

Existing gaps:

Next steps:

### 5. Overall AI Maturity

Define your company's current level of maturity:

☐ Exploring | ☐ Planning | ☐ Implementing | ☐ Scaling | ☐ Realizing

Justify the answer:

---

---

---

---

Explain how far the company is from the next stage:

---

---

---

You can revisit your answers after completing *Chapters 6 to 9*, as you will get new ideas and recommendations to improve the level of progress on all these pillars.

## Sharing your strategic AI roadmap

One of the key aspects of a good AI strategy is to make both the vision and plan accessible and understandable so everyone can know where the organization is going and why certain initiatives are being prioritized (e.g., mandatory training to enable the organization and culture stream, migration to new platforms to improve the technology strategy).

Besides the AI policy from *Chapter 1*, another element that can help you spread the word within the organization is the notion of a strategic AI roadmap. You want to project a long-term vision while keeping things tangible with clear next steps and priorities. For this reason, you can create a multi-year plan (typically between 2 and 4 years) with a "zoom in" for the next calendar or financial year.

The idea behind providing that level of detail for the "Year+1" is to have a series of clear items related to the different AI strategy workstreams. You can see an illustrative example of a 1-year strategic AI roadmap, organized by the five pillars we discussed earlier, as shown in *Figure 5.2*:




Figure 5.2: Example of strategic AI roadmap

Take a moment to review what makes this example effective. Notice how the roadmap starts with the prioritization of use cases and the migration of data systems in the first quarter, because the business and technology foundations condition everything that follows. Notice also how the AI committee review from the internal governance workstream is scheduled before the recurring use cases milestones, so initiatives are validated under clear guardrails before their learnings spread across teams. Finally, the use cases appear as a repeated milestone every quarter, reinforcing the idea that AI maturity is built through continuous, incremental practice rather than one-off initiatives.

The way to consolidate the main priorities, activities, and their expected timing during the yearly planning will depend on who owns the AI strategy and roadmap exercise, but you will typically need to work with the rest of the leadership team and all the management levels to make sure this plan is not an isolated initiative.

There are many ways to do this, but a good one is to establish a calendar with key dates for your road mapping exercise: preliminary planning sessions based on requirements, yearly presentations of the strategic plan, quarterly reviews, etc. Before you move to the next section, spend some time defining not only the potential 1-year roadmap for your organization (aligned with the overall AI vision and strategy), but also the related meetings for the follow-up and review of your initiative over time. Same as with the rest of the exercises, feel free to use your pencil to draft some notes by leveraging the following template:

## Strategic AI Roadmap and Meetings Plan

| 1-Year Strategic Roadmap |    |  |    |  |    |  |    |  |
|--------------------------|----|--|----|--|----|--|----|--|
|                          | Q1 |  | Q2 |  | Q3 |  | Q4 |  |
| Business Strategy        |    |  |    |  |    |  |    |  |
| Organization and Culture |    |  |    |  |    |  |    |  |
| Technology Strategy      |    |  |    |  |    |  |    |  |
| Internal Governance      |    |  |    |  |    |  |    |  |
| Overall AI Maturity      |    |  |    |  |    |  |    |  |
| Main milestones          |    |  |    |  |    |  |    |  |

**Tip**

## Meetings Plan

What are the key AI roadmap-related meetings you need to plan (e.g., design and planning, review)? For each, add the relevant stakeholders from your organization.

Add your comments here:

[illegible]

## Turning your AI strategy into operational actions

By now, you understand the key elements of an AI strategy, along with the typical levels of maturity, and using a strategic AI roadmap as a tool to define and share the key elements of your company-level plans. Here is a last set of recommendations that will facilitate how you convert this AI strategy into reality:

- **Enable a good experimentation-delivery balance:** Generative AI requires continuous experimentation, but value is only realized through disciplined delivery. Establish clear pathways to move from pilots to enterprise-grade solutions, supported by governance, security, and compliance frameworks that scale responsibly.
- **Standardize AI project requirements for sponsor:** Define a common baseline for AI initiatives, including data readiness, infrastructure needs, skill requirements, delivery approach, and ownership. This ensures consistency across projects and shares a clear vision of what the decision factors are for investment and internal executive sponsors.
- **Define and share clear criteria for pilot-production move:** It is a common misconception that every successful experiment should scale; in reality, many pilots exist precisely to test assumptions cheaply and should be discarded if the expected value is not confirmed. Define (or help your technical teams define) standardized decision gates to evaluate technical feasibility, business impact, risk, and scalability before committing additional investment.
- **Define your financial models for AI early:** Clarify when AI initiatives should be treated as CAPEX versus OPEX, and align funding models with expected outcomes, time horizons, and operational ownership. This avoids friction when transitioning from experimentation to sustained operations.
- **Use a top-down strategy to guide bottom-up experimentation:** A clear AI strategy provides direction and guardrails, but you also need to enable teams to experiment autonomously while remaining aligned with enterprise priorities, architectural standards, and ethical guidelines. For that goal, you can:
  - **Block dedicated time for innovation:** Allocate a defined number of hours per week for teams to experiment with AI tools and use cases. This formalizes innovation, signals leadership support, and prevents experimentation from competing with day-to-day delivery work. This is similar to the famous 20% model from Google, in which employees can dedicate one day per week to innovation and upskilling purposes.
  - **Connect ideas to the AI governance or steering committee:** Create structured channels for surfacing, reviewing, and prioritizing AI ideas. The AI committee can



help align initiatives with strategy, manage risk, allocate resources, and accelerate decision-making, so be sure to leverage that.

- **Drive internal AI productivity and adoption:** Focus not only on external-facing AI use cases, but also on internal productivity gains. Provide guidance on where AI adds value, where its limitations lie, and how employees can use AI responsibly and effectively in their daily work. You will see in Chapter 7 how to face internal discussions to make sure everyone understands why this is necessary, and how that impacts their day-to-day activities.
- **Leverage vibe-coding to accelerate experimentation:** Encourage "vibe-coding" as a fast, lightweight way to prototype ideas, explore concepts, and test feasibility, while maintaining clear boundaries between experimental code and production-grade solutions. Ensure that everything aligns with the dos and don'ts of your AI policy, based on the vendor and tool considerations that you will explore in *Chapter 8*.
- **Adopt no-code and low-code platforms to scale impact:** Empower non-technical teams to build and test AI-enabled solutions using no-code and low-code tools. This expands innovation capacity, reduces dependency on engineering teams, and shortens the path from idea to value. As with the previous point, the use of these tools should align with the organizational policy and tooling considerations.

These are just a few recommendations to make sure you are encouraging the internal innovation mindset while respecting the general guidance from your company's overall AI strategy and vision. These best practices can help you accelerate some of the technical, human, and organizational pillars of the AI strategy. If you would like to explore more adoption recommendations of this kind, you can check publicly available resources such as Microsoft's Cloud Adoption Framework, which includes dedicated AI adoption guidance (<https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/>).

Before we conclude this chapter, let me share some recommendations in terms of focus and reading sequence for the next four chapters. You will obviously read all the remaining chapters, but you may choose to adapt the order and start the chapters that are more relevant to your immediate priorities and AI strategy pillars.

## Approaching the next chapters optimally

You can see the reading plans as a strategic lens with justified recommendations based on the current context and needs of your organization. You need to read the scenario and decide which combination looks more suitable for the next stage of AI leadership in your organization. Note that, in the following tables, the chapter references indicate the main theme of each chapter

(business strategy, organization and culture, technology strategy, and internal governance) rather than their full titles. Here are five options (from A to E) you may want to consider:

|                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Option A: Strategy-First (Vision and Value Driven)</b><br><br><i>Chapter 6 → Chapter 7 → Chapter 8 → Chapter 9</i>                                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>When this is the right choice</b><br><br>AI strategy is still <b>being defined or refreshed</b><br><br>Leadership alignment is weak or fragmented<br><br>AI initiatives feel opportunistic rather than intentional<br><br>You are at <b>Exploring → Planning</b> maturity stages<br><br><b>Key benefit of this choice</b><br><br>Prevents <b>technology-led or compliance-led AI</b> , which often fails to deliver business value. | <b>Why this sequence works</b><br><br><b>Chapter 6 (Business Strategy)</b><br><br>Establishes the <i>north star</i> : value pools, competitive advantage, business model impact.<br><br><b>Chapter 7 (Organization and Culture)</b><br><br>Once the "why" is clear, you address the human implications: skills, change, and resistance.<br><br><b>Chapter 8 (Technology Strategy)</b><br><br>Only after understanding the "why", you design platforms and data capabilities that support real priorities.<br><br><b>Chapter 9 (Internal Governance)</b><br><br>Builds governance to enable the strategy, not to constrain it prematurely. |
| <b>Option B: Risk and Compliance-First (Regulated Environments)</b><br><br><i>Chapter 9 → Chapter 6 → Chapter 8 → Chapter 7</i>                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>When this is the right choice</b></p> <p>Your organization is from highly regulated industries (finance, healthcare, public sector)</p> <p>You're facing strong legal, risk, or compliance pressure</p> <p>There are concerns around GenAI misuse, data leakage, or model risk</p> <p>AI initiatives are already happening "in the wild."</p> <p><b>Key benefit of this choice</b></p> <p>Avoids "shadow AI" and retrofitted governance, one of the most expensive mistakes organizations make.</p> | <p><b>Why this sequence works</b></p> <p><b>Chapter 9 (Internal Governance)</b></p> <p>Sets guardrails early, such as acceptable use, risk classification, escalation paths.</p> <p><b>Chapter 6 (Business Strategy)</b></p> <p>Clarifies which AI use cases are worth the risk and investment.</p> <p><b>Chapter 8 (Technology Strategy)</b></p> <p>Selects architecture and vendors that comply with governance constraints.</p> <p><b>Chapter 7 (Organization and Culture)</b></p> <p>Aligns training and change programs with clearly defined rules and expectations.</p> |
| <p><b>Option C: Technology-Led Scale-Up (Platform Consolidation)</b></p> <p><i>Chapter 8 → Chapter 6 → Chapter 9 → Chapter 7</i></p>                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>When this is the right choice</b></p> <p>Strong engineering or data teams are already experimenting</p> <p>There are fragmented tools, platforms, or vendors</p> <p>You're facing difficulty scaling PoCs into production</p> <p>Organization is at <b>Implementing</b> → <b>Scaling</b> stages</p> <p><b>Key benefit of this choice</b></p> <p>Prevents "great models, no impact" by aligning technical scale with business relevance.</p> | <p><b>Why this sequence works</b></p> <p><b>Chapter 8 (Technology Strategy)</b></p> <p>Rationalizes platforms, data pipelines, and integration patterns.</p> <p><b>Chapter 6 (Business Strategy)</b></p> <p>Re-anchors tech investments to value creation and prioritization.</p> <p><b>Chapter 9 (Internal Governance)</b></p> <p>Formalizes controls once scale and impact increase.</p> <p><b>Chapter 7 (Organization and Culture)</b></p> <p>Enables adoption across non-technical teams now that platforms are stable.</p> |
| <p><b>Option D: People and Adoption-First (Change Management Focus)</b></p> <p><i>Chapter 7 → Chapter 6 → Chapter 8 → Chapter 9</i></p>                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>When this is the right choice</b></p> <p>Your organization is seeing high employee resistance or fear around AI</p> <p>There is low AI literacy outside technical teams</p> <p>You're witnessing strong cultural influence on success (e.g., knowledge work, services)</p> <p>You are working with early-stage or stalled AI adoption</p> <p><b>Key benefit of this choice</b></p> <p>Turns AI from a <i>top-down mandate</i> into a <i>bottom-up capability</i>.</p> | <p><b>Why this sequence works</b></p> <p><b>Chapter 7 (Organization and Culture)</b></p> <p>Builds understanding, trust, and baseline capabilities.</p> <p><b>Chapter 6 (Business Strategy)</b></p> <p>Frames AI as a business enabler, not a threat.</p> <p><b>Chapter 8 (Technology Strategy)</b></p> <p>Selects tools with usability and adoption in mind.</p> <p><b>Chapter 9 (Internal Governance)</b></p> <p>Introduces governance that people accept more easily because they already understand the intent.</p> |
| <p><b>Option E: Execution and Maturity-Driven (Fast Value Realization)</b></p> <p><i>Chapter 6 → Chapter 8 → Chapter 9 → Chapter 7</i></p>                                                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

|                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>When this is the right choice</b></p> <p>Executives' pressure for quick, tangible results</p> <p>Clear business use cases are already identified</p> <p>There is a willingness to iterate on culture later</p> <p>Organization is aiming to move from Planning → Implementing</p> <p><b>Key benefit of this choice</b></p> <p>Accelerates ROI while still converging toward a sustainable AI operating model.</p> | <p><b>Why this sequence works</b></p> <p><b>Chapter 6 (Business Strategy)</b></p> <p>Sharpens the prioritization of use cases and KPIs.</p> <p><b>Chapter 8 (Technology Strategy)</b></p> <p>Enables delivery pipelines rapidly.</p> <p><b>Chapter 9 (Internal Governance)</b></p> <p>Introduces governance as scaling begins, not as a blocker.</p> <p><b>Chapter 7 (Organization and Culture)</b></p> <p>Institutionalizes learnings after early wins.</p> |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

There are no better or worse options, but it's best to know that you can play around with the sequence of the next four chapters, based on your unique organizational needs and context. I encourage you to continue reflecting on the current scenario and to decide how you want to proceed at both the AI leadership and reading levels. If you are still unclear, you can keep it simple and follow these recommendations:

- a. What is currently blocking AI progress?
  - Unclear value: Start with *Chapter 6*
  - Fear/Resistance: Start with *Chapter 7*
  - Tech fragmentation: Start with *Chapter 8*
  - Risk exposure: Start with *Chapter 9*
- b. What maturity jump are you trying to make?
  - Strategy-first: Exploring → Planning
  - Tech and execution-first: Planning → Implementing
  - Governance and culture focus: Implementing → Scaling or Realizing

## Summary

In this chapter, we learned how to craft an AI strategy that moves from vision to actual execution. We started by defining the five pillars that support any solid AI strategy: business strategy, organization and culture, technology strategy, internal governance, and AI maturity, along with concrete actions for each of them. We then explored how to assess your organization's AI maturity through the exploring, planning, implementing, scaling, and realizing stages, and you completed a maturity and gap assessment for your own company. We also covered the strategic AI roadmap as a tool to share your vision and priorities across the organization, together with the meetings plan to keep it alive over time. Finally, we reviewed a set of operational actions, from the experimentation-delivery balance to decision gates and financial models, that help convert the strategy into reality and discussed how to sequence the next chapters based on your organization's context.

The next chapter focuses on the business strategy pillar, where you will explore how AI impacts your company's business model, value creation, and competitive positioning.

# 6

## AI-First Approaches: New Business Models and Products

In the early days of enterprise AI, the business justification was fairly straightforward. You could automate decision-making, reduce cost, speed up delivery, etc. Pretty linear impact. And that story still matters, because the idea of automation funds experimentation, proves credibility, and removes friction from workflows. And to be honest, even the new agentic AI patterns are oriented to autonomy and automation, so the justification of impact is technically very similar.

But nowadays, the industry is talking about AI-first approaches, AI native companies, and frontier AI innovation frameworks. The traditional org charts are becoming work charts that combine both human and AI workers. Solo entrepreneurs are creating new startups with just a few tools and agents. AI agents are performing multiple and very long tasks on behalf of the organization. And the economy of this generative AI era is booming and bringing all kinds of opportunities for new products and services. The impact is not linear anymore; it is highly exponential.

So that means there is a growing but still small percentage of companies that are fully embracing AI and bringing it to the center of their solutions. They leverage AI internally for productivity and coding, they consider AI capabilities as a fundamental aspect of the core value proposition they sell to their clients, and they embrace the newest innovations and developments to accelerate and increase their gap versus other competitors.

This chapter is about making that shift in a disciplined way without undue hype, chaos, or breaking trust. It follows the levels of AI maturity that we explored in *Chapter 5* and is based on the understanding that you may be leading a big traditional organization or a newly created startup. You'll learn how AI-first companies are adapting business models and their unit economics, and how to redesign the organization so everyone embraces this mindset and makes the most of the generative AI advantage.



The key topics that we will cover are:

- Understanding the levels of AI transformation
- Rethinking products and services
- The pursuit of scalability
- The AI washing trap
- Managing AI cost units and pricing models

## Understanding the levels of AI transformation

Automation is appealing and tempting. Using automation as your AI leadership mantra is attractive because it is measurable and politically easy to justify: cycle time, error rate, headcount savings, lower budgets, etc. Many organizations start there and then make no progress. They don't reach the fifth "Realizing" phase of the AI maturity model you saw in *Chapter 5* because automation becomes the organization's only mental model for AI. Their technology doesn't evolve because they want proof of value. Their pilots don't evolve because they need a fully confirmed ROI. Every single detail becomes difficult, and what initially was a clear criterion for success ends up becoming a (too) granular and picky exercise that leads to the analysis-paralysis phenomenon, that is, the situation in which teams keep evaluating options, metrics, and risks without ever making a decision or moving to action.

When AI is framed primarily as automation and a cost reduction lever, AI becomes a regular IT project instead of a business transformation, risk avoidance dominates learning and experimentation, and the perception of value is capped by the size of the process being optimized.

One clarification before we dive in: the levels of AI transformation that you are about to explore in this section don't necessarily match the AI maturity stages you saw in *Chapter 5*. You can have a fully new venture that is still implementing or scaling, but making the most of every single piece of AI, every day, and for every part of their product. And you could be in the last Realizing stage while only reaching levels 2 or 3 here. It all depends on the final mix and how AI really transforms your business.

As an AI leader, you can use the idea of automation as a foundation for your business justifications, but you will need to expand into other horizons if you really want your company to go beyond competitors and regular adopters. You can imagine your AI initiative as an area of impact in multiple levels, as shown in *Figure 6.1*:




Figure 6.1 – AI transformation levels

Let's explore each of these levels in more detail:

- **Level 1 – Finding efficiencies:** This level involves the use of quantitative levers that can justify the obvious ROI calculations. The classic activities include reducing cost, accelerating time required for specific tasks, and potentially removing errors in existing workflows. As an AI leader, you will always rely on these areas to justify the short and mid-term financial sustainability of your AI initiative. With the arrival of generative AI, the impact on internal activities becomes highly relevant, as internal work productivity and technical coding tasks are accelerated.  
This level of impact is already happening in a lot of companies, mostly as AI solution adopters that leverage tools like ChatGPT, Claude Code, or any other AI-enabled tool to accelerate and facilitate certain day-to-day activities. The positive aspect is that if the potential efficiency calculation is clear, then the ROI can be more or less certain because this usually requires paying for monthly licenses with a clear budget per user. And it contributes to making some progress in both the technology strategy and the culture and people streams of your strategy. Think, for example, about the many banks, insurers, and software companies that have rolled out coding and productivity assistants to their teams: the business model doesn't change, but the existing work becomes measurably faster and cheaper.
- **Level 2 – Augmenting capabilities:** This is the "crossing the chasm" moment for any AI initiative. Most of the organizations (I would say 80% or 90%, but I don't have the data to back it) stay either as mere AI tool adopters or don't even use AI at scale. But for those moving on and adding AI as a key lever for their business and solutions, AI represents the opportunity to increase the throughput of their decisions via intelligent workflows, automated access to information via copilots, and the introduction of new processes (e.g., automated document analysis and parsing without requiring human intervention, autonomous agents performing some tasks in a sandbox environment).  
The main difference between this level and the first level is that the AI advantage starts to become systematic at the company level, and the aggregated value creation becomes more relevant, not only to automate or reduce cost, but also to accelerate activities that

couldn't be done by hand with the existing human and tool resources. A good real-world illustration is Morgan Stanley, which deployed a generative AI assistant so its financial advisors can query the firm's internal research library in natural language: the point is not only saving time, but augmenting what each advisor can deliver to clients.

- **Level 3 – Transforming business:** This level involves creating totally new offerings, channels, and business models that are not feasible without the technical capabilities of AI. Especially for hyperpersonalization, new access to knowledge, seamless and simplified natural language interactions for atypical non-technical users (e.g., doctors, lawyers), and agentic patterns that help unlock new demand with clients.

Think about the Netflix case. The nature of the transformation wasn't necessarily AI-first, but about moving from traditional physical to digital business models. That said, the way Netflix is reaching users and making sure the recommendations are relevant to them is fully based on AI capabilities (i.e., complex recommender systems).

- **Level 4 – Creating new ventures:** Directly from the third level, sometimes the disruption is so strong that it creates a totally new business. Sometimes it is because of the creation of a new market category (e.g., vibe-coding or quick AI prototyping tools), or because the impact of AI generates a new type of intrapreneurial initiative within the company.

Obviously, this doesn't happen often and certainly not to all organizations. But you can keep this fourth level as something aspirational, and keep in mind that any potential new revenue stream can become a new AI-first firm that evolves from your traditional structure to become a separate entity. A well-known illustration of this dynamic is Waymo, which started as Google's internal self-driving car project and eventually became a separate company under the Alphabet umbrella: a technology-driven internal initiative that matured into a completely new venture with its own business model.

You decide your AI portfolio allocation depending on the funding context and risk appetite. One way to do it if, for example, you still need to prove the value of AI for your organization, is to have a conservative 70% finding efficiencies (to fund the AI engine and prove the ROI), 20% augmenting capabilities (to improve core decisions and raise effectiveness), and 10% to transform your business (new products and services). The last level, if it ever happens, can be funded as a well justified stretch goal beyond the regular portfolio.

#### Tip

Remember, the exact numbers matter less than the system and the discipline, as you want to explicitly manage trade-offs and protect invention capacity from being consumed by urgent automation and hyper-proved ROI work.

A big piece of that transformation idea depends on your ability to evolve your existing products and services, or to create new ones. Let's look at this aspect next.

## Rethinking products and services

The AI component will or can help you improve your existing products, but the key question is: Will this new piece of technology help evolve the existing value proposition of your offering, or will it completely transform that positioning? And how is this different depending on the type of offering?

The transformation can be as simple as a UI evolution in a SaaS platform, or as deep as a truly disruptive way to process the intermediate information within an existing end-to-end workflow. In either case, the transformation is justified by the real value it adds to users. On the other hand, other "AI-first" organizations that position their products and services just based on the idea of AI without justifying its real value and using it as a marketing asset, end up repeating a simple question: What are we actually selling?

In this section, we will explore three complementary angles of this question. First, we will look at the impact of AI on product development and why AI-first products behave differently from traditional ones. Then, we will cover the implementation of AI-enabled services and the new service models that are emerging in the market. Finally, we will review the typical traits of AI-first organizations so you can assess how close your company is to that model.

### The impact on product development

At the product level, remember that traditional products assume users navigate menus and workflows. But AI-first products increasingly assume users express intent in natural language, the system generates content and options, the workflows adapt to context, and the agents can take actions based on permission and controls. Last but not least they include new building blocks and loops related to context (retrieve relevant knowledge + user state), generation and reasoning (draft, plan, recommend), action (execute tasks through tools and systems), feedback loop (capture edits, outcomes, exceptions), and improvements (evaluation, tuning, policy refinement). This is why AI-first products feel different. You are managing a living system, not a static feature set.

While many companies describe their offering as an output or a combination of outputs (data dashboards, policies, reports, diagnoses, lessons, designs), real AI-first companies increasingly define the product as an outcome achieved through an intelligent loop that then generates an output. For example, Harvey or Legora for law intelligence, Google's NotebookLM, or healthcare organizations creating scribes and automated reports for patient-doctor discussions. These examples encapsulate a certain complexity and provide that outcome in a simple, direct way.

This type of reframing matters because it changes the way we implement products. These outputs lead to a series of features, several per outcome. And the new AI outcomes and features are so essential that there would not be a product without the AI piece. We don't see AI as a "feature layer" on top of unchanged workflows or legacy backends. The workflows actually change what users do, instead of just changing what the UI says.

There are already well-known references of this transformation being done well. Think about how GitHub embedded Copilot directly into the developer workflow: the AI doesn't sit in a separate window, it participates in the act of writing code, and the product experience changed accordingly. Or consider how Adobe integrated its Firefly generative models into Photoshop and the rest of Creative Cloud, turning generation into a native part of the editing workflow rather than a separate feature. In both cases, the AI capability changed what users actually do inside the product, which is exactly the difference between an AI-first evolution and a cosmetic add-on.

## Implementing AI-enabled services

The idea is that AI-first services transform not only delivery speed (e.g., accelerating strategy decks or solution development) but also the pricing logic. The notion of human and AI collaboration impacts the way a services organization calculates its cost, margin, and project market value. Some companies even include AI agents in their new headcount (bringing interesting questions about the future of work and the contribution of AI agents to welfare systems).

What is becoming clear today is that many established companies (e.g. integrators, management consulting) are better positioned to win through AI-first services because they already have relationships, trust, domain expertise, and operational workflows. And other new startups are able to leverage AI capabilities to scale their services with just one or a few human employees, reaching new company milestones faster than ever.

This is a relatively new area, but some of the common AI-first service models include:

- **AI-augmented professional services:** Based on the typical types of activities, with faster delivery, higher consistency, and better documentation.
- **Managed AI operations:** In this case, your company runs models, monitoring, governance, and updates for clients who lack maturity.
- **Outcome-based services:** You charge for results, not hours. This means high upside, but also a higher measurement burden and strict risk discipline.
- **Embedded copilots or access to knowledge as a service:** You sell ongoing capability inside customer workflows. Same as other API enabled business models, retention increases with adoption and context depth.

Summarizing, AI-first services often break the traditional "billable hours" equation. If you are working in a services company (we acknowledge that this playbook is more focused on product-led organizations, although most of its guidance can be extended to service-based organizations with some tweaks, as this very section shows), you will need to redesign pricing and delivery models. Otherwise, productivity gains will reduce revenue rather than increase margins or growth. For this matter, you will learn about the key cost units of AI implementations later in this chapter.

## Becoming an AI-first organization

As a consequence of these two product and service models, you can start to identify some of the typical traits of an AI-first company:

- **Building learning loops into operations:** AI-first companies do not rely on occasional "model upgrades." They engineer continuous learning to ensure the following:
  - Feedback capture is intentional (edits, approvals, outcomes, exceptions)
  - Evaluation is systematic (golden datasets, regression tests, red teaming)
  - Improvement is operational (release cadence, monitoring, incident response)
- **Treating "time to innovation" as a core metric:** AI-first companies measure innovation like time spent between steps in a supply chain, from idea to experiment, then pilot, production, and measurable business impact.
- **Using context engineering as a strategic asset:** As generative AI models become widely accessible and increasingly commoditized, differentiation often shifts away from the AI models and toward the context you can add to them. For example, proprietary data and workflows, domain-specific policies and risk rules, customer history and preferences, etc. The context setup becomes one of the core product features.
- **Reusing internal capabilities:** This includes a platform-like mindset to reuse building blocks (prompt patterns, evaluation harnesses, connectors), shared infrastructure (identity, logging, monitoring, policy enforcement), and internal marketplaces (datasets, tools, model endpoints). This enables scale without reinventing everything team by team, as it is part of the internal technology strategy and maturity.
- **Engineering trust as part of the product:** Based on conviction or just a critical need, trust is not a marketing or corporate message, and it becomes engineered and embedded by design. For example, transparency to users ("why this recommendation?"), guardrails (policy checks, safety filters, controlled actions), auditability (logs, versioning, approvals), and human-in-the-loop options. At scale,

trust is not optional for AI-first companies, and it is a competitive advantage (also related to the internal governance pillar of your AI strategy).

- **Designing internal processes to avoid chaos:** The internal development bottleneck is rarely the access to the AI model (assuming that your company has access to the relevant APIs and providers, which is a bit easier in this era of generative AI). The organizational bottleneck is typically process and ownership, but AI-first organizations build a repeatable system for responsible innovation to avoid unclear accountability, slow approvals, fragmented access to data and systems, a lack of evaluation standards, and unclear pathways from prototype to production. AI innovation is expensive, and the time between cycles is critical to stay ahead of the competition, so internal processes are key to avoiding those bottlenecks by standardizing and accelerating innovation.

Let's take some time to analyze this. Obviously, transforming a business model and the related offerings is not trivial. There is no specific way to magically change your company and pretend that it will behave and sell as an AI-first firm. But you can go case by case, analyze the current situation, and see what it will take for you and your firm to go to the next level.

Tip

**AI Transformation and Offering Assessment**

This new assessment combines the company and offering-level considerations from the last two sections. The idea behind it is to understand the current and future impact of AI, not for the company, but for its business model and product/service offering.

Current level of transformation (company-wide)

**Select the estimated level of AI transformation in your organization:**

- ☐ Level 1 - Finding efficiencies
- ☐ Level 2 - Augmenting capabilities
- ☐ Level 3 - Transforming business
- ☐ Level 4 - Creating new ventures

**Main considerations:**

---

---

---

**Key actions to move to the next level:**

---

---

**Impact on the company's offering (product/service)**

Reflect on your current offering. How far are those products and/or services from becoming AI-first? Anything else in the roadmap that would be natively designed as an AI-first offering? What are the key blockers? (technical, conflict with product roadmaps, business model, market readiness, etc.)

**Product/Service #1 \_\_\_\_\_ (add your notes here):**

---

---

---

**Product/Service #2 \_\_\_\_\_ (add your notes here):**

---

---

---

**Product/Service #3 \_\_\_\_\_ (add your notes here):**

---

---

---

**Product/Service #4 \_\_\_\_\_ (add your notes here):**

---

---

---

**Product/Service #5 \_\_\_\_\_ (add your notes here):**

---

---

---

In general terms, you can discuss these topics with the relevant stakeholders from both product and commercial departments. If you happen to lead any of them, you will be uniquely positioned to influence and adjust the new business models for your company and solutions.



These new business models and offerings will only create sustainable value if you manage to scale them beyond the first successful use cases. In the next section, we will explore what AI scalability means for your organization, both from an organizational and a technical perspective.

## The pursuit of scalability

Dying of success shouldn't be an option for you as an AI leader. But it can happen if you don't plan your initiatives properly. Moving from that successful, flagship AI use case that makes your company relevant and helps you reach the "Realizing" stage, to actual enterprise capability requires more than being lucky and getting market traction. And this is a very important issue for some of the top AI-first companies out there.

Revisit the PESTEL framework from *Chapter 2*. Think about the geopolitical dynamics between geographic regions, the scarcity of certain raw materials required for modern chips, or the concentration of power in just a few AI providers. Everything here impacts at a higher level what, by the end of the day, you and your company will find in terms of access to technology.

For example, if resources are limited and your providers cannot offer you certain generative AI models in specific geographies that are key for your activities, or if to do so you need to plan several months in advance to potentially get access to those resources (note: this is the current scenario with plenty of companies working with some of the biggest technology providers), that may affect your ability to scale. And scalability is key to reaching that exponential value from your AI initiative.

The notion of AI scalability includes different aspects. For an AI leader like you, scaling your artificial intelligence investment has at least two dimensions: organizational and technical. Let's discuss these two dimensions in detail.

### Organizational scalability

Organizational scalability represents an operating model shift because traditional organizations tend to fail in one of the following extremes: a centralized AI "ivory tower" that builds prototypes disconnected from business outcomes and ends up failing, or uncontrolled decentralization where every team improvises tools, standards, and risk controls that pose a long-term, cumulative issue.

Even acknowledging the limitations of strategic and tactical leaders to actually influence the operational level, you need to make sure the rules of the game are clear for everyone. Besides sharing your AI strategy (including the vision and roadmap) and policy, you need to clearly

explain the company's organizational choice to help employees understand roles and accountability. There are different models, but some of the classic ones include:

- **Decentralized:** AI capability lives inside individual business units (marketing, ops, product, etc.). Teams build and deploy models tailored to local needs, maximizing speed, domain context, and ownership, but often duplicating tools, data pipelines, and standards.
- **Centralized (a.k.a. Center of Excellence/AI Hub):** A single AI unit owns strategy, talent, platforms, and major models. This improves governance, efficiency, and reuse, but can feel slow or disconnected from frontline needs if demand outpaces capacity.
- **Hybrid or hub-and-spoke:** A small central AI team sets standards, platforms, and reusable models, while embedded AI practitioners sit in business units. Balances scale and speed; success depends on strong coordination and clear decision rights.
- **Federated:** Semi-autonomous AI teams exist in each unit, coordinated by a lightweight central body that aligns on architecture, data standards, and ethics. More flexible than a pure hub-and-spoke, but harder to enforce consistency.
- **Platform-centric:** The AI unit focuses on shared infrastructure (data, MLOps, model access, governance), while business teams build applications on top. Works well in digitally mature organizations.
- **Product-led AI:** AI teams are organized around products or customer journeys (e.g., growth, churn, optimization). Strong accountability for outcomes but can underinvest in foundational capabilities without a coordinating layer.
- **Functional AI (by capability):** Teams specialize by function (e.g., **natural language processing** or NLP, forecasting, computer vision). Encourages expertise and reuse but requires strong intake and prioritization to avoid misalignment with business impact.
- **Innovation lab/Incubator:** A separate AI group explores new use cases, tools, and prototypes with loose constraints. Great for discovery and talent attraction; often struggles to transition work into production without a handoff model.
- **AI as internal consultancy:** A central team is staffed and funded like a consulting group, rotating across units on prioritized problems. Flexible and impact-driven, but knowledge can walk out the door when projects end.
- **Outsourced (partner-led):** Core AI development is handled by vendors or agencies, with a small internal team for oversight. Fast to start and cost-flexible, but risks IP loss, dependency, and weak long-term capability building.

To make these models more tangible, think of the patterns you can observe in the market: large banks and insurers often start with a centralized center of excellence to keep governance tight;

digitally mature technology companies tend to converge on platform-centric setups, exposing shared model access and MLOps capabilities to their product teams; consumer companies with strong product cultures usually prefer product-led or hub-and-spoke approaches; and organizations taking their first steps frequently rely on innovation labs or partner-led models before building internal capability.

If you are wondering which model fits your company, use size and maturity as your first filters. Smaller companies and startups usually get more from decentralized or product-led setups, where speed and ownership matter most. Large organizations at an early maturity stage typically benefit from a centralized model that concentrates scarce talent and sets common standards. As maturity grows, hybrid, federated, or platform-centric models become more attractive because they combine scale with speed. And if AI is not yet strategic for your company, an internal consultancy or a partner-led model can be a pragmatic starting point, as long as you plan how to internalize the knowledge over time.

Regardless of the model you choose, the common thread is the same: clear decision rights, shared standards, and reusable capabilities, combined with enough proximity to the business to stay relevant. Choose the model that matches your current structure and maturity, communicate it explicitly so everyone understands roles and accountability, and be ready to evolve it as your organization scales its AI adoption.

## Technical scalability

You don't need to be an engineer or even a technical leader to set the bar in terms of technical discipline and scalability. As an AI leader, you will need to require these capabilities as non-negotiables for scaled deployments, as a baseline to increase the AI maturity and to improve the technology and human strategy over time. Some of the relevant topics include:

- **Proactive resource planning:** Allocating GPU and other compute capacity in advance to avoid technical bottlenecks and capacity issues from vendors.
- **Evaluation and regression testing:** Every meaningful use case needs an evaluation harness, not just subjective feedback.
- **Observability and monitoring:** Tracking usage, cost, quality signals, and safety incidents in production.
- **Fallback and resilience:** Defining how the system degrades gracefully when uncertain, overloaded, or risk-triggered.
- **Identity and permissions:** Context access must be governed, auditable, and aligned with roles.

- **Release discipline:** Versioning and tracking prompts, tools, retrieval configs, and models to maintain rollback capability.
- **FinOps mindset:** Keeping the cost factor in mind, projecting value creation from cost units, optimizing via reserved instances and cloud commitments.

You don't have to build all of this from scratch. The main cloud providers include evaluation, monitoring, and governance capabilities in their AI platforms (e.g., Azure AI Foundry, Amazon Bedrock, or Google Vertex AI), and there are well-known tools such as MLflow or LangSmith for experiment tracking, evaluation, and observability of LLM-based applications. The point is not the specific tool, but making these capabilities a standard part of how your teams ship AI systems.

To understand why this discipline matters, run a quick, illustrative calculation with your teams. Imagine a product handling 100,000 AI-powered queries per day at an average cost of \$0.003 per query; that is roughly \$9,000 per month. If a model router can send the simplest 60% of those queries to a model that is ten times cheaper, the same workload drops to around \$4,000 per month, without users noticing any difference. Your numbers will obviously be different, but this is the kind of concrete conversation that turns technical scalability into business impact.

In short, the technical dimension is not a detail you can delegate and forget. These non-negotiables are what allow your AI initiative to survive contact with real users, real costs, and real incidents, and they are the foundation on which the organizational model you selected can actually deliver at scale. Answer the following questions to assess your organization's AI scalability:

**Tip****AI Scalability Assessment**

Think in terms of organizational and technical scalability, such as what key items will allow your company to use AI at scale.

**What is stopping (or will stop when the time comes) your company from scaling AI solutions:**

---

---

---

**From the provided organizational items, what are the main strengths of your company?**

---

---

Any promising foundation from the provided technical scalability items?

Create a list of stakeholders and items to address immediate improvements, separating those relevant for the organizational aspect from those relevant for the technical aspect. If you don't know what to improve, feel free to ask them and collect all relevant information here:

## The AI washing trap

**AI washing** is the practice of overstating or misrepresenting the role of AI in a company's products, services, or capabilities in order to appear more innovative. It is one of the most recurrent issues of the generative AI era. Just as many individuals claim their AI expertise and knowledge, there are plenty of companies adding the ".ai" suffix or simply exaggerating the presence of any AI piece to get more market traction and positioning themselves as AI innovators.

This phenomenon has a double impact. On one hand, regardless of how well you work to make your company a proper AI-first firm, you will always deal with "impostors" (understood as those firms that are exaggerating their AI attributes) in the form of competitors and newcomers. On the other hand, as you will see in *Chapter 8*, you will need to pay attention to any external technology provider and evaluate their capabilities to make sure the proposed value proposition is well aligned with their actual capabilities.

But this "AI washing trap" also represents a cycle in which companies become constrained by those claims, leading to disrupted project roadmaps, fragile systems, misallocated investment, governance risk, and long-term erosion of trust.

The reality is that there are no big lessons I can share with you here. Just be careful and avoid this trap. Don't let any market-related efforts accelerate the natural cycle of AI transformation. One approach is collecting feedback from existing and new clients to gather potential interest and details on what they want to see in terms of AI innovation, and a very different (and strongly discouraged) one is to force developments just to feed the market narrative, disrupting the internal operating models along the way. From a context perspective, these are the main stakeholders you want to keep an eye on:

- **Investors:** If you have external investors or a formal executive board, you will certainly need to deal with a proper balance between the industry excitement and the mandate for proper returns. Leverage their great experience and knowledge to understand the market signals, relevant trends, or even competitor actions. But don't hesitate to play your role as the AI leader in your company, and filter potential pressures to adopt AI and to skip your levels of transformation and maturity.  
At the same time, keep in mind that investors want two things: a compelling growth narrative and a credible operational discipline. They will rightfully ask for an AI-specific value proposition, intellectual property, and market advantage, how economics scales over time, what the risks and mitigations are, and what the roadmap is for scaled revenue.
- **Sales teams:** Your sales and business development activities are great sources of feedback from existing and future clients. Working with them will allow you to capture the signals from end-user needs and competitor intelligence perspectives. But keep in mind that in the generative AI era, you will need to transform that information into your own set of AI-first features and offerings (in this case, clients still don't know what is possible or feasible like you do).
- **Product management:** The product teams will be a fundamental partner for you to leverage that useful series of feedback and ideas and to build proper long-term product roadmaps and backlogs. Depending on your leader or C-suite role in the company, you will work more or less close to them. Make sure the feedback loop and product management process are well aligned with the strategic priorities and with your AI vision.
- **Marketing:** Adding the AI flavor to any product or service always looks tempting. And you need to evolve your message and projected value proposition accordingly. But remember that both market and buyers are becoming more mature and aware of the real innovators versus those just falling into the AI trap. Make sure you work with the marketing teams and evaluate relevant assets (e.g. website, one-pagers, product marketing documentation, social media) to make sure they are well aligned with the overall AI initiative.

- **Technical teams:** In this case, the AI trap has a different meaning. AI-related technologies such as models and frameworks, bring excitement and intellectual challenges that motivate your technical teams to explore and learn. And that is great, as your level of maturity heavily relies on their ability to test, experiment, and continue including technical building blocks to their technology stack.
- **Finance/Revenue:** On the other hand, the financial element of your AI initiative will also play an important role. Depending on the context, it can be a powerful ally to create financial and investment plans for your AI efforts, a blocker if they get stuck in level 1 (i.e., looking only for cost reduction and 100% proven ROI), or even a dangerous seed if they themselves fall into the trap and overexaggerate the amount of money available for the AI initiative or the way to spend it, if it brings unbalanced **capital expenditure (CAPEX)** and **operational expenditure (OPEX)** models.
- **Yourself:** Yes, you also need to be careful. The AI washing trap is a comfortable option for everyone. If the rest of the stakeholders are not intentional and you don't explain and manage the risk, the entire initiative will be at risk. That is why this playbook includes so many simple self-assessments, because you need to take time to analyze and reflect on the situation. And you need to do so as an iterative, highly evolving exercise to make sure you capture the 360-degree situation of your organization.

Ultimately, avoiding the AI washing trap is about protecting the credibility of your AI initiative, both externally with the market and internally with your teams. And the best antidote to inflated claims is a clear-eyed understanding of what AI actually costs and how it can be priced. That is exactly what we will cover in the next section.

## Managing AI cost units and pricing models


Generative AI changes cost structures in ways that surprise even experienced leaders. The only way to make sure your company is mastering margin and pricing strategies is by going into the full details of what AI cost represents today, including the main blocks and units, and the different purchasing options.

In this section, we will first break down the AI cost structure into its main dimensions (human, platform, data, and risk and compliance costs) and then dive into the token-based economics of generative AI, with practical recommendations to define sustainable pricing models. These topics are important because they determine whether your AI initiative behaves as a sustainable business or as an ever-growing cost center.

### The AI cost structure

Calculating the aggregated cost for your AI initiatives is a bottom-up exercise in which you and your teams will try to make sense of all the relevant cost units for each of your AI projects. From a

leadership perspective, you need to get the total AI cost to make sure the entire initiative is sustainable and then make sure that every single project is well justified and funded. You can visualize the aggregation of these cost units, as shown in *Figure 6.2*:



*Figure 6.2 – Aggregation of cost units for AI initiatives*

In general terms, the following dimensions are important for you to consider for most of your AI projects:

- **Human costs:** These include everything related to the hourly cost of different roles (technical or not), like the ones we will explore in *Chapter 7*. Depending on your company's human structure, this may include core development, science, and engineering activities, but also workflow design and product management, domain experts for functional evaluation and red teaming, incident response and customer support, change management and training, etc.

One of the key mistakes when planning the human cost of an AI project is to forget about the projected cost beyond the initial development phase, and to focus mostly on the core team working on the technical activities. Of course, the funding mechanism depends on the project context (external or internal), and the sponsor (e.g., paid by a client, funded internally by a centralized entity, or funded by specific departments).

Also, in *Chapter 7*, we will explore the planning and talent considerations related to this human stream, as well as the future of work and impact of AI productivity and coding tools to potentially multiply the team's bandwidth with a relatively low investment (e.g., AI copilots, on-demand cloud capabilities).

- **Platform costs:** It is a category that contains any cost related to tools, licenses, first or third-party infrastructure, and, in the case of AI projects, the access to models such as LLMs from companies like OpenAI or Anthropic.



From the infrastructure perspective, there are some baseline costs related to general computing power, servers, and other core capabilities such as storage, networking, and observability. They are not exclusive to the AI portion of your developments, but these elements impact the overall availability, reliability, and latency of your systems. As you already know, you can get access to these capabilities by setting up your own physical infrastructure or by renting these capabilities.

Besides those baseline costs, there will be another variable aspect related to training, evaluation, and inference of AI models. You will learn more in the next section about the economics of generative AI and how you can make sense of it to understand future investments.

From the tooling perspective, and depending on the type of tool (e.g., productivity, coding, knowledge management), you will be paying a monthly fee for individual or team licenses to access specific capabilities that can accelerate your AI initiatives and multiply the overall development and project velocity. Those are easy to quantify as they usually follow linear patterns, such as *number of licenses  $\times$  monthly price*, or they involve master agreements with technology vendors to get access to privileged commercial conditions.

- **Data costs:** It is often an overlooked category of AI project costs. Think about it for a bit. You can get access to data in very different ways. You can:
  - Own it, if your company collects any relevant information from its products and services
  - Acquire it, if you pay a data license to another company that may have relevant proprietary information
  - Have free access to it from public statistics and open data portals
  - Generate it, for example, with synthetic data generation techniques to expand existing datasets that may be too small for your AI model training purposes

But regardless of where the data comes from and the related economics, there are multiple tasks (that translate into human and tooling costs too) that will be required to collect, clean, adapt, aggregate, analyze, and prepare the information to be AI-ready.

This stream often connects with company-level data governance and platform initiatives in which the data sources are already being managed, including some relevant mechanisms for data privacy and compliance, access control and security, and lineage (i.e., understanding the nature of the data transformations from the origin). Leveraging

this, depending on the overall data maturity of the organization, may not only simplify your projects but also lower the final data cost per project.

- **Risk and compliance costs:** You can see **risk and compliance costs** as a double stream. On one hand, you can think about the cost related to implementing proactively all sorts of controls that will enable privacy and audits, compliance with regulatory requirements, legal reviews for intellectual property, risk mitigation and liability, technical AI safety testing, etc. On the other hand, you can calculate the cost of inaction, which could be way higher as it would imply mandatory audits, reactive approaches, paying fines, and dealing with potential reputational damage.

As you will see in *Chapter 9*, this cost stream will include technical, compliance, and organizational activities that will not only handle risk and compliance but also enable responsible AI approaches for your organization. This will make everything easier and better for both the firm and you as an AI leader.

When you sponsor the AI projects, or you are at least part of the decision-making process to decide which AI projects should be accepted, besides keeping in mind the overall AI strategy, vision, and maturity goals, ask for the numbers.

If you are part of the executive or leadership team, you are probably not the person who will prepare those numbers, but you can help the team members in charge (e.g., project managers, product owners, delivery managers) validate the hypotheses and calculations.

Note that the weight of these cost dimensions is closely related to the levels of AI transformation we discussed at the beginning of this chapter. At the first level, costs are dominated by licenses and subscriptions, which are relatively predictable and behave like classic operational expenses. As you move toward augmenting capabilities and transforming the business, platform, data, and human costs grow and become more variable, and risk and compliance costs deserve dedicated attention. And if you ever reach the point of creating new ventures, you are essentially funding a startup, with the full cost structure and investment logic that this implies. Keep this progression in mind when you project budgets for the next stages of your

AI journey.

## The token economy in the era of generative AI

Pricing is where generative AI initiatives become sustainable businesses, or expensive and unrealistic experiments. The key item that will impact your cost estimation and, therefore, your potential margin and revenue models is one of the technical terms you have seen in previous chapters: the tokenization of the information.

Tokens are the units into which a model breaks the content it processes (roughly, fragments of words, and in multimodal models also pieces of images or audio), as introduced in *Chapter 4*. Providers bill by input and output tokens, which makes them a fundamental cost driver.

When you think in terms of tokens, you are trying to translate regular user-tool interactions such as clicks or natural language queries, or internal processes that automatically consume multiple documents containing X words, and to translate those into their equivalent of Y tokens. When something happens with your platforms and that action involves an AI capability, then you will need to cover the cost of those AI model calls, generally through an API connection that you pay to technology companies like Anthropic, Microsoft, OpenAI, Google, AWS, Meta, etc.

But deciding your own pricing based on this type of cost is hard. Generative AI introduces variability because cost depends on factors like input/output length, model type and version, provider, context depth and memory, and the complexity of the processes (e.g. multi-step agent workflows). From a vendor perspective, even if they don't explain the units behind their token-based billing, token usage often correlates with the compute and cost required to make models run (plus their corresponding margin).

All this complexity means that calculating cost and margins, and therefore pricing, cannot be separated from product design and operational discipline. Here are a few recommendations:

- **Pick the unit you will manage:** Decide early what you are actually optimizing for: cost per customer, per workflow, per transaction, per outcome, or per active user. This choice shapes everything from UX to architecture. If you don't define the unit upfront, pricing drifts, usage explodes in unexpected places, and margin surprises appear after scale, not before. The "wrong" unit often leads teams to optimize the wrong behaviors (e.g., more usage instead of more value).
- **Select the right setup intentionally:** Final quality and performance perception varies with cost, latency, and risk controls. Higher-end models may improve accuracy or UX but introduce latency or cost volatility; cheaper setups may degrade trust or outcomes. Be explicit about where you are willing to trade cost for quality (or vice versa), and align that choice with user expectations, SLAs, and brand positioning.
- **Develop a FinOps mindset:** AI-first leaders build **token literacy** across product, engineering, and finance, similar to how cloud-native companies build **cloud cost literacy**. **Financial operations (FinOps)** teams should see, understand, and discuss generative AI costs in near-real time. *Free* usage creates waste and a lack of accountability at scale, so introduce internal chargeback or showback mechanisms. This allows you to allocate costs granularly by feature, product area, customer tier, or workflow and make better prioritization decisions. The end goal is to have a common

understanding of the involved cost and to decide a target margin that guides all decisions.

- **Decide your pricing pattern:** There is no single correct pricing model. You can start from hypotheses and usage scenarios, then set pricing based on maximum thresholds, average usage, or average cost per user tier. Some products bundle AI costs into plans, others gate usage, and some rebill AI consumption directly while adding a margin per token or per request. The key is to choose a pattern that aligns incentives between you and your users and keeps your downside risk under control.
- **Account for total cost of ownership (TCO):** Model cost is only one part of the equation. When choosing between cloud and local environments, proprietary and open-source models, factor in infrastructure, engineering effort, maintenance, observability, security, compliance, vendor lock-in, and upgrade cycles. A cheaper model per token can be more expensive overall if it increases operational complexity or slows iteration.
- **Control environments and usage:** Put guardrails in place: quotas, per-user or per-workspace limits, usage tiers by customer type, and rate limits for expensive workflows. These controls protect margins, prevent abuse, and make costs predictable, while also giving you levers to upsell or differentiate premium tiers.
- **Optimize at the technical level (when ready):** Early AI projects prioritize speed and learning, not efficiency. But as teams mature, technical optimization becomes a competitive advantage. Explore building blocks like model routers that automatically select the most suitable model based on complexity, risk, latency needs, or customer tier. Invest in retrieval design, caching, summarization, and memory policies. These techniques can significantly reduce token usage, improve margins, and even allow you to lower prices while delivering better performance.

Instead of doing another exercise or self-assessment, my honest recommendation is to sit with your teams and have discussions. Discuss products, AI features, cost hypotheses, and try to understand how your organization is making sense of the business model, product management, branding, cost, and pricing topics related to your AI initiative. You can add your notes here:

**Tip**

**Final notes for Chapter 6 topics:**

---

---

---

---

[illegible]

## Summary

In this chapter, you learned how to move beyond automation and turn AI into a true business transformation lever. We started by exploring the four levels of AI transformation, from finding efficiencies and augmenting capabilities to transforming the business and creating new ventures, and how to balance your portfolio across them. We then discussed how to rethink products and services with an AI-first mindset, including the impact on product development, the emerging AI-enabled service models, and the typical traits of AI-first organizations. We also covered the pursuit of scalability from both organizational and technical perspectives, the AI washing trap and the stakeholders you need to watch, and, finally, the AI cost structure, the token economy, and practical recommendations for pricing your AI-enabled offerings. With these tools, you can now assess where your company stands and define a credible path toward becoming an AI-first organization.

In the next chapter, we will focus on the human side of this transformation: the talent, roles, and team structures you will need, and how AI is reshaping the future of work itself.

# Part 3

## Building the Organization for AI at Scale

*In Parts 1 and 2, you laid the foundations for your AI leadership approach and your company's transformation. But AI becomes real through people, technology choices, and responsible execution. Part 3 focuses on building and reskilling AI-ready teams, navigating an increasingly complex ecosystem of tools and vendors, and establishing the ethics, governance, and accountability required to scale AI sustainably. Together, these chapters help leaders build the organizational capabilities and guardrails that turn AI ambition into lasting execution.*

This part contains the following chapters:

- *Chapter 7, Building and Scaling AI Teams: Roles, Reskilling, and Hiring*
- *Chapter 8, Navigating the AI Ecosystem: Tools and Vendor Evaluation*
- *Chapter 9, Responsible AI Implementation: Ethics, Governance, and Accountability*



# 7

## Building and Scaling AI Teams: Roles, Reskilling, and Hiring

None of the topics and areas of knowledge matter if you don't have a proper human structure in your organization. You cannot translate AI strategy into actionable plans, and you cannot create a company culture that makes the most of AI technologies to release new products and services. And even if you set the tone and enable an AI-first culture, the reality is that you need the right people and skills to make it happen.

Obviously, AI talent is relatively scarce and difficult to secure due to fierce competition among companies. And you will need to figure out a way to attract new hires and retain your top talent by combining the benefits package with other factors such as your company's mission, the impact of their work, the time to learn and experiment, etc. When we talk about AI talent, we mean professionals who combine some of the following skills: data and software engineering, model development and evaluation, prompt and context design, AI product management, domain expertise, and AI governance and compliance. The most reliable way to evaluate these skills is against the actual responsibilities of the role, using realistic scenarios and measurable outcomes rather than titles or certifications alone. We will explore the specific roles and their skillsets in the *Understanding the evolution of roles in the AI development cycle* section later in this chapter.

For this and other reasons, your role as a key AI leader will be critical to building and scaling AI teams in your organization. Let's see how. We'll be covering the following topics:

- Your own role as an AI leader
- Understanding the evolution of roles in the AI development cycle
- Other considerations for your AI talent strategy



- Key considerations to scale your AI teams
- Defining the operating model

## Your own role as an AI leader

Even if the job market is highly disrupted by restrictive immigration policies, economics, and technology disruptions (including automation and AI), the reality is that AI talent (not only technical people, but also product, design, analytics, legal, and domain professionals working on AI initiatives, as we will see later in this chapter) is still uniquely positioned and has enough skills and levers to allow everyone working in this industry to thrive and choose where and how they want to work.

As with any other competitive activity, it will require more than just salary and benefits to convince the best professionals to join your company instead of going somewhere else, including your competitors.

From the AI leadership perspective, you can differentiate your company by being part of the talent process, including hiring and retention. The steps below combine talent attraction with culture, retention, and ethical responsibility, since these dimensions are inseparable in any sustainable AI talent strategy. Let's look at some steps you can take.

## Own the AI narrative

You will need to personally shape and communicate how AI fits into the company's mission and operating model. During the hiring interviews or even all-hands and 1-on-1 meetings with employees, you can own that narrative by doing the following:

- Clearly articulate why the company is investing in AI and what problems it exists to solve, beyond generic productivity gains.
- Translate the AI strategy for different internal audiences (engineering, product, sales, legal, leadership), not just technical teams.
- Be explicit about ethical boundaries, governance, and responsible AI use; top talent increasingly evaluates companies on these dimensions.
- Ensure AI teams understand how their work connects to real business impact and decision-making authority, not just experimentation.

This way, you can help talent feel more aligned with the company's AI vision and purpose, instead of just following and completing tasks. It also helps connect the strategic plan to their day-to-day activities.

## **Establish visible AI thought leadership**

Remember that high-caliber professionals want to work with leaders who have opinions, not just roadmaps or numbers. You can do this by:

- Developing and sharing a clear point of view on AI trends, trade-offs, and future directions (even when uncertainty exists).
- Encouraging leaders to publish, speak, or contribute to discussions on AI strategy, deployment, and organizational change, not only technical breakthroughs.
- Internally, hosting regular forums or talks where leadership discusses AI bets, failures, and learnings openly.
- Signaling curiosity and intellectual honesty, as strong talent is usually attracted to leaders who evolve their thinking publicly.

As a consequence of this, the company becomes a learning platform, not just a workplace. This builds trust with your current teams and makes the company a more appealing option for future candidates.

## **Design pragmatic AI interview processes**

Hiring is one of the strongest signals of how seriously you take talent, and the first impression your best candidates will get about your company. Getting a final match requires a lot of serendipity, and you need to make sure your company is doing its best to make the interview process as useful and positive as possible. You can do the following:

- Make interviews pragmatic and grounded in real-world scenarios the candidate would actually face. Exaggerating the interviews with unrealistic scenarios will only make your company lose the best talent. Limit the number of interviews, schedule them relatively close to each other (for the candidate to keep the focus on the process), and don't ever propose free, unpaid x-day trials.
- Avoid the excess of administrative processes. There are plenty of opportunities for the candidates to drop the process if they need to interact with multiple recruiters, HR, compensation calculation, and offer teams (all this added on top of the regular interviews). Senior and top talent evaluates the company just as closely as the company evaluates them.
- Involve future teammates and cross-functional partners to reflect how work actually happens, but once again limit the number of interviews and clarify the process upfront.
- Focus on judgment, prioritization, and communication, instead of just technical depth or academic knowledge. And when your teams evaluate the technical component, make

sure they know how to interview and guide the candidates so the interview doesn't feel like an adversarial, mind reading exercise for them.

- Be transparent about the expectations, and don't look for perfection. AI talent will likely bring a variety of cultures, languages, and personalities. Focus on the core value of those candidates; otherwise, you will miss great, atypical talent.

By the end of the day, you will make better hiring decisions and develop stronger trust with newcomers from day one. Keep in mind that proper hiring and onboarding processes are key to retain any top talent, including AI.

## **Enable feedback loops for the AI initiative**

Retention is driven by whether people feel heard and see progress. And your AI strategy won't really translate into proper actions if your talent is not convinced or involved in the decision making. Remember to:

- Establish structured channels for employees to share ideas at both the company and project level (e.g., quarterly reviews, innovation forums, anonymous feedback). Focus on the AI strategy pillars, related activities, priority projects, etc.
- Actively close the loop. If you embrace feedback, you are in charge of communicating what was decided, done, or adopted, what wasn't, and why. Treat and leverage feedback as a leadership input, not an HR exercise (that will only create an excess of process, and a bottleneck to really make the most of the direct employee feedback).
- Encourage feedback not just on tools and models, but on processes, incentives, and cross-team collaboration. This is especially important for continuous improvement, AI management systems, AI strategy pillars, etc.

With so many moving pieces in your AI initiative, you really need to master this part. The best AI talent feels and embraces ownership, but not just in exchange for employment.

## **Tell a credible long-term story**

Remember that top talent optimizes for trajectory, not just compensation. The best paid AI professionals get a generous mix of high salaries, stock, social benefits, etc. But, except if they are engaged in a one-shot kind of consulting mandate, they give value to multi-year plans with high potential. Regardless of your level of financial resources, try to adapt the approach and ensure the following:

- Articulate a multi-year vision for how AI capabilities will evolve inside the organization. Align the AI story across talent, investment, and client narratives to avoid misalignment.
- Show how today's hires and current employees can grow into tomorrow's leaders as the AI footprint expands. That scalability is key to show potential to grow. And remember

that it is not always about becoming people managers or executives; AI talent can evolve into senior, principal, lead, and distinguished kinds of roles over time.

So, remember that people join for the future, not just the role. That future includes vertical growth, but also lateral movement across teams, domains, and projects, which many AI professionals value as much as promotions. Focus on both company- and talent level paths to success and establish a system that works and helps you retain and promote talent over time.

## **Analyze the market and understand the investment**

Design remuneration levels, career paths, and growth mechanisms aligned with the reality of AI work. You will need to:

- Work with the rest of the leadership team, HR, and the hiring managers to understand the market. This includes not only salaries, but also job descriptions and detailed skills. It is important to know what you are getting into, and that knowledge will apply to both newcomers and existing employees. Understanding the market also means understanding your specific domain: skills, expectations, and compensation for AI roles can vary significantly between industries such as financial services, healthcare, retail, or technology.
- If it fits your company's policy, one of the most welcomed best practices in the AI industry is to show the salary range in a transparent way. Top AI talent won't waste time learning about the salaries later, and that will often be a key decision factor for them to join or not the interview series.
- Be realistic. Same as with the salary and the rest of the benefits, you can list your requirements, and it is fair, but be careful. Don't ask for 10 years of LLM experience because it doesn't exist, and evaluate first if you can ask for top senior talent with very extensive experience if you cannot really afford it. Showing high expectations and low salaries is often a way to send the best talent far from your hiring reach.
- Make expectations for advancement explicit, including required skills, scope of responsibility, and measurable impact at each level. Revisit and adjust career paths regularly as AI technologies, tooling, and organizational needs evolve, ensuring roles do not stagnate or become misaligned with market expectations.

These are key aspects to guarantee a frictionless process, but also a regular baseline for any hiring initiative. Make sure these foundations are ready when you try to find external talent for your AI projects.

## Empower your AI teams

Enable your teams to reach their maximum potential, and create psychological safety and meaningful decision autonomy for AI teams. This sounds like a classic ABC for any role, but it is especially important for AI talent. Remember to:

- Establish cultural norms that encourage experimentation, constructive disagreement, and honest discussion of model limitations, data quality issues, and technical trade-offs.
- Protect teams from blame when well-reasoned experiments fail, and treat failures as learning inputs rather than performance liabilities. You can also delegate real decision-making authority on tools, architectures, and implementation choices, signaling trust and reducing friction caused by excessive top-down control.
- Plan for knowledge management and retention. Leverage the knowledge transfer mechanisms as a way to reduce risk in a very volatile job market where rotation is high and talent leaves very often.

These are good practices that you can support from the executive level and replicate in the multiple levels of AI management, so the employee experience is similar regardless of the team, department, or geography.

All these practices define how you attract, retain, and empower AI talent as a leader. To apply them effectively, though, you also need a clear picture of who that talent actually is, which roles you will be hiring for, and how those roles are changing. Let's explore this next.

## Understanding the evolution of roles in the AI development cycle

For you to know how to manage your AI talent, you need to understand the variety and nuances of the multiple roles involved in the AI development cycle. Even more, you will need to make sense of the evolution of these roles, often related to industry trends and technology disruptions like generative AI.

The following diagram summarizes this landscape, showing the main types of AI-related roles and how traditional profiles are evolving into new, generative AI-era positions, as shown in *Figure 7.1*:




Figure 7.1 – Types and evolution of AI-related roles

Let's look at some of the types of roles you will hire and manage, depending on your organization's structure and talent priorities.

## Traditional roles

Traditional software and data roles are not going anywhere. Even in the generative AI era, they remain the backbone of any AI initiative: they build, secure, and operate the systems that AI features run on, and most of the newer roles we will discuss later are evolutions of these profiles. Let's review them one by one, focusing on how each contributes to your AI strategy.

## Frontend Software Developer

A frontend software developer focuses on building user interfaces and client-side experiences for web and mobile applications, translating product requirements and design systems into responsive, accessible, and performant UI components that integrate with APIs and backend

services. In an AI context, they increasingly embed AI-powered features such as chat interfaces, copilots, recommendations, and dynamic content generation into user workflows while ensuring usability, latency management, and transparency of AI outputs.

From an AI leader's perspective, frontend developers are critical to responsible AI adoption: they shape how AI is perceived by users, implement visual guardrails in the interface (e.g., disclaimers, feedback loops, confidence indicators), and collect structured interaction data for model improvement. AI leaders must ensure frontend teams understand AI limitations, human-in-the-loop design, and experimentation frameworks to safely and effectively deploy AI features at scale.

## **Backend Software Developer**

Backend developers design and maintain server-side logic, APIs, databases, integrations, and system performance, ensuring reliability, scalability, and security of applications. They also handle authentication, rate limiting, logging, and cost optimization for model usage. From an AI leader's standpoint, backend developers are foundational to operationalizing AI: they control how models are invoked, how data flows securely, and how output is validated before reaching users. Leaders must align backend architecture with governance policies, observability standards (model latency, drift, hallucination rates), and cost controls, ensuring AI capabilities are production-grade rather than experimental prototypes.

## **Cloud Administrator and Engineer**

Cloud experts manage infrastructure provisioning, deployments, networking, security configurations, and performance optimization across platforms like AWS, Azure, or GCP. In AI environments, they oversee GPU provisioning, model hosting environments, containerization (e.g., Kubernetes), serverless orchestration, and hybrid or multi-cloud architectures. They ensure compliance with enterprise security and data residency requirements.

Cloud engineers determine whether AI initiatives can scale sustainably: leaders must collaborate with them to balance innovation with cost, resilience, and compliance. Strategic considerations include vendor lock-in risks, data governance in cloud AI services, and establishing standardized MLOps/LLMOps pipelines that support experimentation without compromising enterprise-grade stability.

## **Data Engineer and Architect**

Data experts design data pipelines, ingestion frameworks, storage systems, and data models that ensure clean, reliable, and accessible data for analytics and AI workloads. They build pipelines for document ingestion, embedding generation, metadata enrichment, and real-time streaming to support retrieval-augmented generation and personalization. They also enforce data quality, lineage, and governance controls. Data engineers are central to AI performance: model quality is constrained by data quality. You will need to prioritize structured data governance, privacy-by-

design principles, and scalable vector infrastructure, ensuring that proprietary knowledge assets can be safely leveraged while maintaining compliance and minimizing bias.

## **Data Scientist and ML Engineer**

Data scientists and ML engineers develop, train, fine-tune, evaluate, and deploy AI models, applying statistical analysis and experimentation to solve predictive or optimization problems. They design evaluation benchmarks and manage model lifecycle processes, including monitoring drift and retraining. They bridge experimentation and production, often collaborating closely with MLOps teams. Responsible experimentation frameworks, reproducibility, and model risk assessments become key governance priorities, so these roles are very important for your company's success.

## **Quality Assurance and Security Engineer**

QA and security engineers test software systems for functionality, robustness, vulnerabilities, and compliance, ensuring applications are resilient against misuse or attack. In AI-driven systems, their role expands to include adversarial testing, prompt injection defenses, red teaming, bias testing, and validation of AI outputs under edge cases. They assess data leakage risks and enforce secure model integration practices. For AI leaders, QA and security are strategic safeguards: generative systems introduce novel risks (hallucinations, jailbreaks, sensitive data exposure). Leaders must embed AI-specific testing protocols into **software development lifecycle (SDLC)** processes and fund ongoing monitoring, recognizing that AI risk management is continuous rather than one-time certification.

## **UX Designer**

UX designers research user needs, design workflows, create prototypes, and test interactions to ensure products are intuitive, accessible, and aligned with user goals. In AI-powered products, UX expands into conversation design, trust calibration, explainability, and collaborative human-AI interaction models. Designers must anticipate uncertainty in AI outputs and craft experiences that communicate system boundaries clearly. UX determines adoption and trust: even the most advanced model fails if users don't understand or trust it. Leaders should prioritize AI literacy within design teams and champion user testing specifically focused on AI behaviors, transparency cues, and human override mechanisms. They often collaborate with front-end engineers to convert designs into implemented interfaces.

## **Product Analyst/Data Analyst**

Product and data analysts interpret user behavior, performance metrics, A/B tests, and funnel analytics to inform product decisions. In AI-enabled systems, they measure AI feature usage, output quality signals, cost per interaction, conversion uplift, and user satisfaction metrics tied to AI assistance. They help define KPIs that reflect business impact rather than novelty. Analysts will provide you with the evidence base for scaling or pivoting AI initiatives: decisions about



expanding model usage, fine-tuning strategies, or retiring features should be data-driven. Leaders must ensure analytics frameworks capture AI-specific metrics such as hallucination rates, fallback usage, and human override frequency.

## Other tactical roles

Roles such as Project and Product Managers, Product Owners, Scrum Masters, Technologists, and Strategists coordinate product vision, backlog prioritization, stakeholder alignment, delivery processes, and technology roadmaps. In AI initiatives, they translate ambiguous AI capabilities into concrete use cases, define MVPs, manage experimentation cycles, and align technical feasibility with market opportunity. They also manage cross-functional dependencies among data, engineering, compliance, and design teams. From an AI leader's perspective, these roles operationalize strategy: you must ensure AI roadmaps align with core business value, risk tolerance, and long-term capability building rather than opportunistic feature additions. Clear governance, phased rollouts, and measurable milestones are essential.

With the arrival of generative AI, some of these roles evolved and expanded their area or scope of activities. We'll look at these roles next.

## New roles

For some of the previous roles, their day-to-day changed, and even their formal titles. Most of these new positions did not appear out of nowhere: they build on the foundations of the traditional roles we just reviewed, adding generative AI-specific capabilities such as prompt and context design, model orchestration, and AI-specific evaluation and governance. Let's look at some of the new trends and their evolution from their traditional counterparts.

### Generative AI Developer

They specialize in building applications that integrate foundation models via APIs or hosted services, implementing prompt orchestration, retrieval mechanisms, tool use, and workflow automation. They understand token limits, temperature tuning, cost trade-offs, and context engineering. From an AI leadership standpoint, GenAI developers are the accelerators of innovation, rapidly prototyping new AI-enabled experiences. Leaders must ensure they operate within architectural and governance guardrails, with shared prompt libraries, evaluation standards, and monitoring frameworks to prevent fragmented or inconsistent AI implementations across teams.

### Generative AI Cloud Expert

This role combines deep expertise in cloud infrastructure with specialization in AI-native services such as managed LLM platforms, vector databases, model hosting, and GPU optimization. They evaluate and configure services like Microsoft Foundry and Azure OpenAI, Amazon Bedrock, Google Vertex AI, or Hugging Face deployments. From an AI leader's perspective, this role is

strategic in vendor selection, cost forecasting, scalability planning, and risk mitigation. Leaders must rely on them to assess trade-offs between managed services and self-hosted models, ensure data isolation, and design architectures that support both experimentation and enterprise resilience.

## **Generative AI Architecture**

They design end-to-end AI-native systems, integrating data pipelines, model layers, retrieval systems, orchestration frameworks, security controls, and user interfaces into cohesive, scalable solutions. They define patterns for RAG, agent-based systems, multi-model strategies, and hybrid cloud deployments. From an AI leader's viewpoint, this role is pivotal in preventing architectural sprawl: architects translate strategy into reference architectures and standards that guide multiple teams. Leaders should empower them to define reusable frameworks, enforce governance standards, and anticipate future scaling needs, ensuring AI investments compound rather than fragment.

## **Prompt Engineering**

Prompt engineers design, test, refine, and automate prompts to optimize LLM performance for specific tasks, often building structured templates, evaluation datasets, and guardrails. They experiment with system instructions, few-shot examples, and output formatting to achieve consistency and reliability. While the standalone title may evolve, the skill remains crucial in early-stage GenAI deployments. From an AI leadership perspective, prompt engineering should mature into systematic prompt management and evaluation practices integrated into CI/CD pipelines. Leaders must ensure that prompt strategies are documented, version-controlled, and aligned with compliance requirements rather than treated as ad hoc experimentation. Last but not least, this kind of profile rarely stays a standalone position: it tends to converge with broader profiles.

Many prompt engineers evolve into general AI engineers, who combine prompting with software integration and evaluation work, or into context engineers, a more recent term for specialists who design everything the model receives, including data, tools, instructions, and memory, beyond the prompt itself. You can consider both as natural evolutions of the prompt engineering skill set rather than entirely separate roles to hire for.

## **Generative AI PM and others**

Generative AI-focused product managers and related tactical roles specialize in identifying high-impact AI use cases, defining AI feature requirements, and balancing innovation with user trust and compliance. They work across legal, security, engineering, and business units to launch AI responsibly. From an AI leader's standpoint, these roles bridge executive vision and execution, translating strategic AI ambitions into prioritized roadmaps with measurable ROI. Leaders should ensure these PMs understand both technical constraints and ethical implications,

enabling them to make trade-offs that protect brand reputation while driving competitive advantage.

## AI Regulation and Responsible AI Expert

AI regulation and **Responsible AI (RAI)** experts monitor evolving laws, standards, and ethical frameworks governing AI and generative use, such as data privacy regulations, AI-specific legislation, fairness mandates, and transparency requirements. They develop internal policies for model governance, bias mitigation, explainability, and auditability, often collaborating with legal and compliance teams. From an AI leader's perspective, this role is not merely defensive but strategic: proactive compliance builds trust with customers, regulators, and partners. Leaders must integrate RAI considerations into product design and architecture decisions early, ensuring AI systems are not only innovative and profitable but also lawful, ethical, and socially responsible. With the arrival of generative AI, this is even more important due to the increasing level of AI adoption and the specific requirements from regulations such as the EU AI Act or the Colorado AI Act in the USA, among others.

## Forward Deployed Engineer

One of the most recent roles is the **Forward Deployed Engineer (FDE)**. Very common in the most advanced AI and technology companies, the FDE operates at the intersection of engineering, product, and customer engagement, embedding directly with clients to implement, customize, and scale complex technical solutions in real-world environments. In a Generative AI context, FDEs configure model integrations, adapt prompts and workflows to client-specific data, implement retrieval pipelines, fine-tune orchestration layers, and troubleshoot performance or alignment issues in production settings. They translate abstract AI capabilities into tangible business outcomes by rapidly iterating on deployments, integrating with legacy systems, and navigating organizational constraints such as compliance, security, and change management.

From your perspective as an AI leader, FDEs are strategic accelerators of adoption and revenue: they shorten time-to-value, surface real-world edge cases that inform product and model improvements, and act as a feedback bridge between customers and core engineering teams. You can imagine FDEs not merely as implementers, but as intelligence nodes in the ecosystem, capturing insights about scalability, usability, regulatory friction, and ROI while ensuring they operate within standardized architectural, governance, and responsible AI guardrails to prevent bespoke deployments from fragmenting the broader platform strategy.

## Other critical roles

Beyond the extensive set of AI-related roles that we've just discussed, there are others that are critical to scaling AI responsibly and effectively across an organization. Some of these are the following:

- **Subject Matter Experts (SMEs)** provide deep domain knowledge that grounds AI systems in real operational context, ensuring use cases are relevant, outputs are accurate, and edge cases reflect real-world complexity rather than theoretical assumptions; they are essential for reducing hallucinations and aligning models with business nuance.
- **AI Committees** (often cross-functional governance bodies, as you saw in *Chapter 1*) and **AI Centers of Excellence (CoE)** are relevant too. The former typically bring together executives, legal and compliance leads, security officers, and senior technical profiles such as AI architects; the latter usually combine AI architects, engineers, data experts, and trainers. For example, the AI CoE can lead standardization, best practices, reusable architectures, and capability building across business units, preventing fragmented experimentation and ensuring shared learning compounds over time.
- **Trainers** (internal or external) accelerate adoption by upskilling teams, embedding AI literacy, promoting responsible usage, and translating complex technical concepts into practical workflows.

Remember, all these roles collectively determine whether AI remains a series of pilots or becomes an institutional capability embedded in culture, economics, and governance. But it doesn't mean that you need to hire all of them. Try to adapt your roles and skillsets to the context, resources, and level of maturity of your organization.

## Other considerations for your AI talent strategy

In this section, we'll look at some additional considerations that impact the talent exercise and the way the industry uses their skills for AI-related initiatives.

### Shifting from science to engineering

Over the past decade, the *Data Scientist* moved from being labeled the "sexiest job of the 21st century" to becoming a more normalized (still very important but not always critical, depending on the type of AI project) role within AI-driven organizations. Initially, scarcity of statistical and ML expertise, combined with hype around big data, elevated data scientists as strategic differentiators. However, as tooling matured (AutoML, foundation models, no-code analytics, pre-trained APIs) and generative AI abstracted much of the modeling complexity, the bottleneck

shifted from model experimentation to productionization, integration, scalability, and reliability (all of which are domains dominated by engineering).

Today, value creation increasingly depends on shipping AI systems, not just building models. This has repositioned data scientists from isolated model builders to collaborative contributors within cross-functional AI product teams. From an AI leadership perspective, the key evolution is not a decline in importance but a reframing of impact: the highest-leverage data scientists now focus on evaluation frameworks, experimentation design, model governance, and translating ambiguous business problems into measurable AI objectives, while partnering tightly with engineers who own deployment and infrastructure. Organizations that fail to rebalance toward strong engineering foundations often accumulate impressive notebooks but limited business outcomes. The modern hierarchy of value is less about prestige and more about proximity to scalable execution.

## Selecting between fundamental and applied research

This is a classic nuance that you need to understand properly before allocating talent to your AI initiative. Fundamental research involves advancing core knowledge for developing new algorithms, improving model architectures, conducting theoretical research, exploring novel training methods, and publishing or contributing to the scientific frontier. These activities push the boundaries of what is technically possible but may not immediately translate into commercial value.

On the other hand, applied research activities focus on using existing models and techniques to solve concrete business problems, integrating APIs, optimizing prompts, building retrieval pipelines, ensuring compliance, measuring ROI, and improving user workflows.

While fundamental research drives long-term differentiation for frontier labs and hyperscalers, most enterprises generate value through execution speed, domain adaptation, integration quality, and operational reliability. Confusion between these domains often leads to misallocated resources: companies attempt fundamental research without the talent density or scale to compete, when their competitive advantage lies in domain data, distribution, or operational know-how. Strategic clarity requires defining whether the organization aims to invent new capabilities or extract value from existing ones. Both are valid, but they demand different talent profiles, investment horizons, risk tolerances, and performance metrics.

## Avoiding trends and focusing on skills and responsibilities

In rapidly evolving AI markets, organizations often hire based on trending titles rather than clearly defined responsibilities, leading to misalignment, underperformance, and internal confusion. FDEs are a prime example (similar to prompt engineers in the first wave of generative AI, a few years ago): some companies recruit them assuming they are simply "AI engineers for clients," others treat them as sales engineers, implementation consultants, or hybrid product

managers. In reality, FDEs are execution specialists who embed with customers to operationalize complex solutions, translate business needs into technical implementations, and feed real-world insights back into product and engineering teams. Without clarity, companies risk creating roles that overlap with solutions architects, customer success, or core engineering, diluting accountability.

In your case, discipline in organizational design is essential: define problems and skills first, and then roles. Define ownership boundaries, tie responsibilities to measurable outcomes, and avoid copying titles from high-profile tech companies without understanding their operating model.

Before moving on, take some minutes to complete a short assessment and analyze the current and future scenarios for your talent initiatives. Think in terms of roles and skills, and both internal (upskilling, reskilling) and external (hiring) levels.

The maturity scale used in the following assessment box goes from initial planning to full consolidation:

- **Very initial:** You are still planning your first AI hires. AI work, if it happens at all, is done ad hoc by existing employees, without formally defined AI roles or evaluation criteria.
- **Increasing:** You already have dedicated AI talent, but it is mostly concentrated in a single team. The rest of the organization depends on that team for anything AI-related, which can quickly become a bottleneck.
- **Expanding:** New talent is joining different areas and projects across the organization. Roles are becoming more specialized, although practices, standards, and career paths are not yet homogeneous.
- **Consolidating:** You are hiring externally and developing talent internally at the same time, with defined roles, career paths, and reskilling programs. AI skills are treated as an organizational capability, not as the attribute of a single team.

Let's get started with your assessment:

#### Tip

##### AI Talent Assessment

Analyze your existing talent structure. Take some time to understand how well aligned it is to your AI ambitions and current level of maturity.

Level of talent maturity (well defined roles, variety of skills):

- ☐ 1) Very initial: Still planning first hires
- ☐ 2) Increasing: Mostly allocated to one single team
- ☐ 3) Expanding: New talent joining different areas and projects

☐ 4) Consolidating: Hiring and developing new talent internally

Key existing (AI-related) roles:

---

---

---

Next potential hires (for external talent):

---

---

---

Feasible upskilling (for existing non-AI team members):

---

---

---

## Key considerations to scale your AI teams

In this section, we'll discuss some hot topics that you will need to proactively address if you want to build and scale your AI teams while maintaining sustainable growth and relationships with your team members.

### Analyzing talent gaps and planning

A rigorous AI talent strategy begins with mapping current capabilities against the organization's 2–3 year AI ambition, distinguishing between experimentation needs and scaled production requirements. This includes inventorying existing skills (data literacy, software engineering maturity, domain expertise, change management capacity), identifying structural bottlenecks (e.g., lack of MLOps, weak data governance, no AI product ownership), and assessing role clarity versus title inflation.

You need to segment talent gaps into build, buy, borrow, and automate decisions, determining which capabilities must be developed internally for strategic differentiation (e.g., AI architecture, governance), which can be hired externally (e.g., specialized ML expertise), which can be accessed via partners, and which can be abstracted through tooling. Workforce planning must also account for velocity. AI adoption compresses timelines, so sequencing matters. For example, in an ideal world, governance and architecture foundations should precede aggressive AI scaling.

The outcome should be a phased capability roadmap aligned to business priorities, not a reactive hiring spree driven by market hype.

## Reskilling the current workforce

Effective AI reskilling operates on three complementary layers:

- **Strategic reskilling** targets executives and senior leaders, building fluency in AI economics, risk management, operating model shifts, and portfolio thinking so they can make informed prioritization decisions rather than delegating blindly to technical teams.
- **Tactical reskilling** focuses on managers, product owners, analysts, and operational leads, enabling them to identify AI-suitable workflows, redesign processes, define measurable success metrics, and manage hybrid human-AI teams.
- **Technical reskilling** equips engineers, analysts, and domain specialists with hands-on capabilities such as prompt design, workflow automation, data structuring, API integration, and evaluation methods.

The objective is not to turn everyone into ML researchers, but to create layered AI literacy where decision-making authority matches understanding. From a leadership standpoint, reskilling is a capital allocation decision: investing in internal capability often yields higher long-term ROI and cultural alignment than over-reliance on external hiring.

## Addressing the future of work issues

This means clarifying vision and impact, explaining areas of automation, and creating psychological safety for both blue and white collar professionals. Your AI transformation must be accompanied by an explicit narrative about how work will evolve. Leaders should articulate a clear vision distinguishing **augmentation** (AI assisting humans), **automation** (AI replacing specific tasks), and **innovation** (AI enabling new roles and services). Transparency about which activities are likely to be automated (e.g., routine documentation, repetitive reporting, standardized coding tasks, transactional service interactions) reduces rumor-driven anxiety and low morale.

Equally important is explaining where human judgment, creativity, relationship-building, and accountability remain irreplaceable. For blue-collar roles, this may mean augmentation through predictive maintenance, robotics guidance, or safety analytics. For white-collar roles, it may mean AI agents for drafting, analysis, and decision support. Psychological safety emerges when employees understand that adaptation pathways exist: reskilling programs, role transitions, internal mobility, and clear performance expectations in AI-augmented environments. Without this clarity, your productivity gains will be impacted by disengagement and resistance.



## Ignoring the noise and avoiding the AI layoff trap

You have seen this everywhere. Reacting to AI hype with premature layoffs or sweeping cost-cutting narratives erodes trust, damages employer brand, and often undermines long-term capability. Declaring automation-driven headcount reductions before systems are even stable creates fear-driven behaviors that weaken the AI transformation efforts.

Instead, you need to justify hiring and workforce plans through capability logic: where AI increases leverage, complementary skills (integration engineers, governance leads, domain translators) may actually increase in importance. Transparent communication about workforce evolution builds credibility externally and internally. At the same time, organizations must prepare for the rise of "build your own solution" behavior, where employees use tools like OpenAI Codex or Claude Code to recreate or customize internal tools. Rather than suppressing this as shadow IT, forward-looking leaders establish safe sandboxes, governance APIs, reusable components, and security guidelines that channel grassroots innovation into sanctioned pathways. The strategic objective is controlled empowerment: enabling bottom-up productivity gains while maintaining architectural coherence, compliance, and data protection.

The considerations we just covered tell you what to address as your AI teams grow: talent gaps, reskilling, the future of work conversation, and communication discipline. The missing piece is how to organize all this talent so it can actually deliver. That is the purpose of your operating model, which we will define next.

## Defining the operating model

The core organizational challenge in AI transformation is internal alignment and coordination. AI cuts across traditional boundaries: product defines the use case, engineering integrates it into workflows, data teams prepare and manage inputs, legal and compliance interpret regulatory exposure, security manages data risk, and operations must sustain the system over time. When these groups only intersect at late-stage approval gates, delivery slows dramatically. Friction accumulates in handoffs, rework increases, and accountability becomes diffuse.

AI-first organizations recognize that speed and safety are not opposites. Instead of serial approvals, they embed cross-functional ownership directly into delivery teams so that risk, engineering, and domain expertise evolve together rather than sequentially.

When it comes to structuring AI teams, most organizations choose between three classic models. In a centralized model, all AI talent sits in a single unit that serves the entire organization; this maximizes knowledge sharing and standards, but can become a bottleneck and lose touch with business domains. In an embedded model, AI professionals are distributed across business units and product teams; this maximizes domain proximity and speed, but risks duplicated effort, inconsistent practices, and weaker governance. A hybrid model combines a central enablement

group with distributed delivery teams, aiming to capture the benefits of both. The structure described next is a practical implementation of the hybrid model.

An effective AI operating model often follows a *Pods + Platform + Embedded Governance* structure, in which cross-functional pods deliver business outcomes on top of a shared platform, with governance built into both, as shown in *Figure 7.2*:




Figure 7.2 – The Pods + Platform + Embedded Governance operating model

Let's look at these different components in detail:

At the front line are **AI Product Pods**, aligned to business domains such as sales, customer service, underwriting, or supply chain.

- A typical pod includes some of the roles you have seen before. For example, a product leader responsible for outcomes, adoption, and value capture, engineers responsible for workflow integration, system reliability, model management, evaluation, etc., domain experts who ensure quality and handle edge cases, and others. Everything together and at the same time. Each pod owns measurable business outcomes within its domain and operates with end-to-end accountability. This structure obviously introduces complexity, but it also increases alignment and prevents "throw it over the wall" dynamics and ensures that risk and feasibility shape design early.

- You can also have a centralized **AI Platform Team**, which provides shared enablement rather than owning business outcomes. The platform team manages the model gateway and routing logic, evaluation harnesses and standardized test suites, retrieval infrastructure with identity and permission controls, monitoring and logging systems, incident response tooling, and approved privacy and safety patterns. It also produces internal documentation, reusable templates, and training resources. By centralizing infrastructure and guardrails while decentralizing outcome ownership, the organization balances autonomy with coherence.
- Governance in this model becomes a capability embedded in tools, standards, and workflows, not only a committee that reviews artifacts after decisions have already been made.

If you choose this kind of operating model, you need to make ownership and decision rights explicit. For example, pods own outcomes (value delivered, adoption rates, and workflow performance within their domain) and the platform team owns enablement (reusable tooling, technical standards, architectural integrity, and system reliability). Risk and compliance functions own policy definition and evidence standards, while partnering directly inside delivery teams to operationalize them. Executives and managers like you own prioritization and trade-offs, including portfolio allocation, sequencing, and risk appetite. When these boundaries blur, AI efforts either stall or fragment. Clarity of ownership is required to avoid politics and chaos.

Last but not least, remember that incentive design determines behavior. If leadership measures only cost reduction, teams will prioritize automation at the expense of quality or innovation. If they measure only the number of experiments shipped, they risk fragmentation and technical debt. If they measure only model performance metrics, they may achieve impressive benchmarks without meaningful adoption. AI-first organizations like yours need to balance metrics across several dimensions: measurable business value, adoption and behavior change, quality and safety standards, and cost-to-serve discipline, and implementation velocity (the time it takes to move from idea to validated outcome). This balanced model ensures that AI initiatives generate sustainable impact rather than short-term optics.

The same dimensions give you a practical scorecard for AI team performance and cross-functional collaboration: for example, adoption of pod deliverables by business users, reuse of platform components across teams, cycle time from idea to validated outcome, and the number of initiatives delivered jointly by technical and non-technical roles.

To help assess these dimensions systematically, the following *AI Operating Model Assessment* provides a practical framework for evaluating your organization's readiness and performance:

Tip

**AI Operating Model Assessment**

Reflect on the topics you just learned about. Try to understand your current or potential structure.

**Existing or future pod structure (based on roles from previous assessment):**

---

---

---

---

**Platform team details and maturity:**

---

---

---

---

**Incentives and bonuses for AI-related topics:**

---

---

---

---

**Current risks and deficits, and potential improvements:**

---

---

---

---

---

## Summary

This chapter was an initial exploration to help you define your human structure based on skills and responsibilities, rather than just roles and trends. We started with your own role as an AI leader: owning the AI narrative, establishing visible thought leadership, designing pragmatic interview processes, enabling feedback loops, telling a credible long-term story, understanding the market, and empowering your teams. We then reviewed the evolution of AI-related roles, from traditional software and data profiles to generative AI-era positions such as the generative AI architect and the forward deployed engineer. We also discussed the shift from science to engineering, the choice between fundamental and applied research, talent gap analysis, the three layers of reskilling, the future of work conversation, and the *Pods + Platform + Embedded Governance* operating model. Keep in mind that the human element is as important as your technology stack: take care of your talent (while still asking for excellence), and they will take care of your AI initiative.

In the next chapter, we'll explore approaches for managing your relationships with technology vendors and internal tool adoption, which are key elements for your AI strategy and to continue increasing the level of maturity of your teams and initiative.

# 8

## Navigating the AI Ecosystem: Tools and Vendor Evaluation

You have already reviewed several aspects of your AI strategy and leadership approach. While defining proper AI-first business models and managing AI talent will help you gain traction and maturity, there is one aspect that can be either a key enabler or a real blocker for your AI initiative: the mix of tools and platforms, and the relationship with your technology vendors.

As a leader and decision maker in your company, you are already feeling the commercial heat. Depending on the context, you may get solicited and unsolicited proposals from multiple technology companies, claiming the unique value of a myriad of tools and solutions.

Besides their individual value propositions, there is an undeniable reality. Commercial goals and tactics often outweigh the realistic technical capabilities of those tools, and there is a lot of noise between competitors sharing a mix of justified and biased statements to undermine the rest of the companies. For example, during the first wave of generative AI, with the release of ChatGPT, many companies accelerated their commercial and competitive statements even before having anything to offer. The era of half-truths, financial-year cycles, and individual bonus motivations is here, as it was before, but it is now more difficult to make sense of it due to the changing and accelerating nature of AI technologies.

This gap between commercial promises and technical reality is now well documented. Market analysts have long described it through frameworks such as Gartner's Hype Cycle, which shows how emerging technologies reach a peak of inflated expectations before the market corrects itself. More concretely, regulators have started to act on so-called AI washing, the practice of overstating the role or maturity of AI in a product. In 2024, for example, the U.S. Securities and Exchange Commission publicly sanctioned investment firms for making misleading claims about their use of AI (see the SEC's enforcement press releases at <https://www.sec.gov>), and the U.S.

Federal Trade Commission launched enforcement actions against companies for deceptive AI claims (<https://www.ftc.gov>). The lesson for you as a buyer is simple: treat vendor claims as marketing hypotheses to be verified, not as technical facts.

This chapter will help you cut through that noise. Its goal is to give you a structured, repeatable way to understand the AI technology landscape, evaluate tools and vendors against your own needs, and design a resilient technology strategy that keeps you in control of the relationship with your providers. You will learn how the different layers of the AI technology stack fit together, how to define selection criteria and run a practical vendor evaluation exercise, how cloud technologies impact your AI strategy and financial models, and how to preserve flexibility while avoiding vendor lock-in. These learnings matter because your tooling and vendor decisions will shape your costs, your talent needs, your compliance posture, and, ultimately, your ability to deliver AI value at scale.

Before we separate real value from noise, let's first examine the layers of the AI infrastructure stack so we can distinguish the different types of tools available.

The following key topics are covered in the chapter:

- Understanding the AI technology stack
- Evaluating tools and vendors
- Leveraging cloud technologies for AI strategy
- Building resilience while avoiding vendor lock-in

## Understanding the AI technology stack

The notion of technology stack refers to the set of tools that coexist and interact to enable your AI developments. In general terms, you can see it as layers that start from the bottom and keep adding new layers to complement your technology environment with additional capabilities. For example, physical infrastructure that enables platforms, platforms that help your teams accelerate AI development, studios and frameworks that facilitate that development and turn capability into products, and governance that ensures sustainability and trust. Together, they form a vertically integrated AI capability that goes from raw compute to end product deployments.

To make this more tangible, think of how these layers materialize with a specific provider. In the Microsoft ecosystem, for example, Azure datacenters with GPU-enabled compute represent the infrastructure layer; managed services such as Azure Machine Learning and Azure AI services provide the platform layer; Visual Studio and playgrounds allow teams to prototype, test, and evaluate models; open-source frameworks such as PyTorch or LangChain support the actual development work; and security, monitoring, and responsible AI tooling provide the governance

layer on top. Equivalent combinations exist for AWS, Google Cloud, and other providers. *Figure 8.1* illustrates these layers:



*Figure 8.1 – AI technology stack layers*

Each of these layers has a specific function and includes a variety of technology pieces, which we will describe here, even if you don't necessarily need to know all the granular details.



## AI infrastructure

At the base of the stack, AI infrastructure is the foundational compute, storage, networking, and data backbone that powers model training and real-time inference, determining performance ceilings, scalability, latency, and the unit economics of AI across the enterprise. This layer determines what is technically and economically possible, and every layer above it depends on the capacity and economics defined here. It can include the following components:

- **Compute foundation:** GPUs/TPUs, high-performance CPUs, networking, and storage required to train and serve models
- **Cloud and on-prem environments:** AWS, Azure, GCP, hybrid setups that determine scalability, latency, and cost structure
- **Model hosting and inference infrastructure:** Containerization (Docker), orchestration (Kubernetes), model servers
- **Data infrastructure:** Data lakes, warehouses, streaming systems, feature stores.
- **Scalability and reliability mechanisms:** Autoscaling, load balancing, observability, fault tolerance
- **Cost management levers:** Compute optimization, workload scheduling, model compression

## Platforms

Sitting directly on top of the infrastructure layer, platforms are the managed environments and orchestration layers that abstract infrastructure complexity, standardize workflows, and enable repeatable development, deployment, integration, and monitoring of AI systems at scale. This layer abstracts infrastructure complexity and accelerates development velocity. Pay special attention to this, and it will impact your company's ability to implement and deliver value, and it will also guide the type of talent and skills you will need to acquire. This layer includes the following:

- **Model and agent platforms:** Managed AI services (e.g., foundation model providers, ML platforms)
- **MLOps/LLMOps platforms:** Experiment tracking, CI/CD for models, monitoring, retraining pipelines
- **Data platforms:** Data labeling systems, data integration and orchestration, governance tooling
- **Integration platforms:** APIs, workflow automation, enterprise system connectors.
- **Security and access management systems:** Identity, role-based access, secrets management

## Development studios

One level up, development studios are the collaborative workspaces where teams design, prototype, test, and refine AI solutions, turning raw models and capabilities into usable applications through experimentation, iteration, and validation. This is where the initial ideas will become actual, working AI products. This layer includes:

- **Application development environments:** IDEs, prompt engineering workspaces, notebook environments
- **Model experimentation spaces:** Sandboxes for prototyping, fine-tuning, and evaluation
- **Testing environments:** Red-teaming, bias testing, robustness validation
- **Collaboration tooling:** Version control, model registries, shared prompt libraries
- **Rapid iteration frameworks:** Low-code/no-code AI builders, copilots

## Libraries and frameworks

Cutting across the development process, libraries and frameworks are the reusable technical building blocks that accelerate development by providing standardized tools, model architectures, orchestration patterns, and optimization methods to construct intelligent systems efficiently and reliably. These choices are usually driven by your technical teams, but they matter to leaders because they influence hiring needs and interoperability with the platform layer. Some of the blocks include the following:

- **Core ML frameworks:** PyTorch, TensorFlow, JAX
- **LLM-related frameworks:** LangChain, LlamaIndex
- **Agent interoperability protocols:** MCP, A2A
- **Model optimization libraries:** Quantization, distillation, acceleration toolkits
- **Evaluation and monitoring libraries:** Metrics frameworks, explainability toolkits
- **Domain-specific libraries:** Specific code libraries for computer vision, NLP, speech, recommendation engines

## Governance, compliance, and safety

Finally, governance, compliance, and safety form the oversight layer that embeds accountability, risk management, regulatory alignment, and ethical safeguards into the AI lifecycle, ensuring systems remain secure, transparent, compliant, and trusted over time. This layer ensures AI systems are trustworthy, lawful, and aligned with organizational values. Rather than sitting on

top of a single layer, it is a transversal function that touches all the previous ones. You will explore this in *Chapter 9*, but here are some of the main blocks:

- **Model governance frameworks:** Version control, documentation, audit trails
- **Risk management processes:** Bias detection, fairness assessments, impact analysis
- **Regulatory compliance mechanisms:** GDPR, AI Act readiness, sector-specific regulations
- **Security controls:** Data privacy safeguards, adversarial testing, model protection
- **Ethical AI oversight:** Transparency standards, human-in-the-loop review
- **Operational monitoring:** Drift detection, performance degradation alerts, misuse detection

To see how these layers work together in practice, consider a common scenario: an enterprise chatbot that answers employee questions using **retrieval-augmented generation (RAG)**. The application runs on cloud infrastructure that provides the compute capacity and hosts the document indexes; an AI platform manages the model endpoints, data pipelines, access control, and monitoring; the team builds and evaluates the assistant in a development studio, iterating on prompts and testing answer quality; frameworks such as LangChain or LlamaIndex orchestrate the interaction between the model and the company's knowledge bases; and the governance layer defines what data the chatbot can access, how its answers are logged and evaluated, and how the system complies with internal policies and external regulations. Every layer plays a role, even in a seemingly simple use case.

From a vendor perspective, you may work with one single company that aggregates most of these capabilities in one single offering, or with multiple companies and tools that are interoperable based on standards and interfaces such as APIs and direct integrations. Let's now see how to structure your analysis and evaluation process to optimize your technology choices.

## Evaluating tools and vendors

Navigating the complexity of building blocks in your technology stack, along with the multiple companies and vendors involved, can be overwhelming for any leader. Sometimes the decisions need to come from the top if they involve commercial negotiations and master agreements. But other times it will be an organic bottom-up choice from your technical teams, often choosing specific coding languages and frameworks, and asking for funding if their needs involve financial investment.

## Defining the AI tooling selection criteria

Part of your role as an AI leader in your company is to establish the tooling choice criteria, which has both technical and procurement considerations. This means choosing the best solutions for

the specific needs and context of the organization, but also aligning the evaluation process to the official buying and procurement considerations of your company.

There are a few factors that you may want to consider when choosing AI tools that can become part of your own selection checklist:

- **Vendor profile:** Evaluate the vendor's background, including their industry expertise, years in business, and financial stability. Choosing an AI tool is also betting on their ability to deliver new innovations in the short, mid, and long term. Here's the criteria to look for:
  - **Reputation:** Research customer reviews, case studies, and analyst reports to gauge the vendor's reliability, trustworthiness, and past performance. Sources like Gartner or Forrester, among others, are usually a good high-level indicator.
  - **Pace of innovation:** Assess how frequently the vendor updates and improves their AI tools, ensuring they stay competitive with industry advancements. Their industry partnerships with other AI companies may help you understand how well positioned they really are.
  - **Geographic considerations:** For example, if you are located in Europe, you may want to understand the considerations of working with tool vendors from other geographies, such as the United States, Canada, or China. This may have implications in terms of applicable regulations, data transfer, sovereignty, etc.
- **Tool scope:** This may include determining whether the AI tool covers a broad range of functions or is specialized for a specific task within your business needs, such as the following:
  - **Environment:** Understand if the tool is hosted in the cloud or as an on-premise option. Both have different advantages, and if you choose a cloud-based tool for AI, you also need to understand the level of service (infra IaaS, platform PaaS, software SaaS, or model MaaS). You can explore these definitions separately in the following note if you have never heard about them, as this will impact your **total cost of ownership (TCO)**. From a cost perspective, the two environments behave very differently. Cloud-based tools require little or no upfront investment and are billed as you use them, which makes them ideal to start quickly and scale progressively, although sustained heavy usage can make recurring costs significant over time. On-premises options require considerable upfront investment in hardware (especially scarce and quickly depreciating GPUs), facilities, and specialized staff, but they can become cost-competitive for organizations with stable, high-volume workloads, strict data residency needs, or existing datacenter capacity. Your TCO analysis should compare both scenarios

over a multi-year horizon, including hardware refresh cycles, energy, and people costs, and not just the license or subscription price.

- **Cloud service levels at a glance: Infrastructure-as-a-service (IaaS)** gives you access to raw compute, storage, and networking that your teams configure and manage, offering maximum control with maximum operational responsibility. **Platform-as-a-service (PaaS)** provides managed environments (e.g., machine learning platforms or managed databases) where the provider operates the underlying infrastructure and your teams focus on building. **Software-as-a-service (SaaS)** delivers ready-to-use applications on a subscription basis, with no infrastructure to manage. Finally, **model-as-a-service (MaaS)** exposes AI models through pay-per-use APIs, so you can consume state-of-the-art models without hosting them yourself.
  - **Features:** Decide if the tool provides only raw AI compute power or includes additional capabilities like data preprocessing, model training, governance, compliance, and workflow automation.
  - **Type of models:** Check if the tool supports different AI models (ML, DL, generative AI, etc.) that fit your business requirements. You may want to analyze the type of model license to understand if they leverage proprietary or/and other open-source models. This matters in the real world because model licenses differ widely in what they allow: some open-weight model licenses restrict commercial use, redistribution, or usage above certain scale thresholds (popular 'open' models often ship under custom, non-OSI licenses), while proprietary terms may include clauses about data usage and intellectual property. If you skip this check, you may face legal exposure, unexpected licensing costs, or a forced migration if a license changes or your usage falls outside the permitted terms.
  - **AI knowledge scope:** Choose between general-purpose AI models and industry-specific AI solutions tailored for particular industries and use cases. It really depends on how you are planning to use the tools; there is no right or wrong choice here.
  - **Compliance and security:** Ensure the AI tool meets industry compliance standards (e.g., EU AI Act, CCPA, Loi 25, GDPR, HIPAA) and includes robust security measures to protect data. In practice, you can verify this by requesting the vendor's certifications and audit reports (e.g., ISO/IEC 27001, SOC 2, or the newer ISO/IEC 42001 for AI management systems), reviewing their data processing

agreements and data residency options, and involving your legal, privacy, and security teams in the evaluation before signing.

- **Usage considerations:** Analyze the internal user profile, as well as how tool pricing/cost scales with usage, ensuring that the tool can support your growth without becoming prohibitively expensive.
  - **Technical level:** Consider whether the tool requires deep technical expertise or if it supports no-code/low-code interfaces for broader accessibility within your team.
  - **Interoperability:** Verify upfront whether the AI tool integrates smoothly with your existing tech stack, including cloud platforms, databases, and third-party software. This could be via API integrations or specific industry connectors.
  - **Customization:** Evaluate how flexible the tool is in terms of customization, API integrations, and adapting to your specific business workflows.
  - **Service level:** Assess the vendor's support structure, including response times, dedicated account management, and availability of technical assistance.
  - **Pricing model:** Understand whether the AI tool vendor will charge you based on usage (pay-as-you-go) or a fixed subscription model, and how this aligns with your budget. For both cases, make sure you get a proper estimation based on expected volume and growth scenario, including potential calculator tools directly created by the vendor.

All these factors are relevant, so you can review them as a checklist. Let's start by choosing one of the tools from your technology stack and doing an initial exercise:

#### Note

##### Individual AI Tool Review Checklist

Choose one tool and vendor and add your initial thoughts here.

General information

**Vendor:** \_\_\_\_\_

**Tool name:** \_\_\_\_\_

**General description:** \_\_\_\_\_

**Pricing model:** \_\_\_\_\_

**Internal considerations:** How well does this AI tool align with our company's strategic goals, existing workflows, and team capabilities?

\_\_\_\_\_

**Contractual aspects:** Are there any hidden costs, vendor lock-in risks, or unclear billing structures that could pose challenges down the line?

**Cost vs value tradeoff:** Does the potential value of this AI tool (in efficiency, automation, or innovation) justify its total cost of ownership?

**Potential ROI:** What measurable returns in terms of cost savings, productivity gains, or revenue growth can we realistically expect from adopting this AI tool?

**Alignment with internal AI usage policy:** Does this AI tool comply with our organization's ethical standards, governance policies, and risk management frameworks?

This checklist is intentionally presented as a blank worksheet so you can make it your own, but you will not be starting from zero: the illustrative exercise that follows, and the worked example at the end of this section, show how these fields look once completed for a specific use case and vendor.

Let's now do an illustrative exercise that you can later adapt to your own company context.

**Illustrative exercise:** ACME Advanced Manufacturing (AM) Corp is a leader in advanced manufacturing, specializing in precision engineering, additive manufacturing, and automation

solutions for aerospace, automotive, and medical industries. Headquartered in New York, the company operates multiple production facilities across the Americas and maintains strategic partnerships with global suppliers and research institutions.

With over two decades of experience, ACME AM has established itself as a key industry player. The company leverages cutting-edge technologies such as 3D printing and robotic assembly to optimize production efficiency and reduce material waste. ACME AM adheres to strict regulatory standards, including ISO certifications, to ensure compliance with industry best practices and quality assurance protocols.

ACME AM serves a diverse client base, ranging from aerospace suppliers to automotive **original equipment manufacturers (OEMs)** and medical device manufacturers. Its competitive advantage lies in its ability to produce complex, lightweight components with high structural integrity, meeting stringent performance requirements.

**Current needs and challenges:** As ACME AM expands its operations and integrates Industry 4.0 technologies, the company faces several challenges that necessitate the adoption of **artificial intelligence (AI)** solutions. Let's number these challenges (1-4) to keep track of each need later in this exercise.

1. **Process optimization and predictive maintenance:** The company seeks AI-driven analytics to enhance real-time monitoring of manufacturing equipment, predicting potential failures before they occur. By implementing AI-powered predictive maintenance, ACME AM aims to minimize downtime, extend machinery lifespan, and optimize resource allocation.
2. **Supply chain and inventory management:** AI-driven demand forecasting is crucial to improving inventory control, reducing excess stock, and mitigating supply chain disruptions. The company seeks AI solutions to enhance procurement strategies by analyzing market trends, supplier reliability, and logistics constraints.
3. **Quality control and defect detection:** Traditional quality inspection methods are time-intensive and prone to human error. AI-based computer vision systems can automate defect detection and ensure compliance with rigorous industry standards. Also, ML algorithms can refine quality assessment models, improving accuracy and reducing material wastage.
4. **Customization and scalability:** The company aims to implement AI-driven generative design tools to optimize product development and accelerate prototyping. AI-powered simulation models will enable rapid adaptation to shifting customer demands while maintaining cost efficiency.

While the value of AI for ACME AM is clear, they don't know how to get started. Keep in mind these four challenges for later, as your practical exercise is about to start.



**Exercise:** You will be the interim AI leader of ACME AM. You are helping the company frame the AI needs before making any tool selection decision. You are getting preliminary, partial information about potential use cases and tools. By the end of the exercise, you will leverage several templates to help you complete your analysis. Let's get started.

Firstly, along with the rest of the team, you have prepared several ideas for use cases:

- a. AI-powered predictive maintenance system
- b. AI-driven energy optimization
- c. AI-based demand forecasting
- d. AI-powered supplier risk assessment
- e. Computer vision-based defect detection
- f. AI-enhanced quality predictive analytics
- g. Generative design for product optimization
- h. AI-driven mass customization platform
- i. AI-powered digital twin for rapid prototyping

This was part of a very preliminary ideation. You don't have detailed information, but you know what the use case would be, you imagine its potential value, and you can anticipate the areas of the company that would make more use of it. You are starting to think about how they are related to the four challenges we have previously mentioned.

Additionally, you have received several vendor tool proposals as part of a **request for information (RFI)**, with three AI vendors answering the call:

- A. **Cloud-based SaaS for text-to-blueprint 3D draft generation:** A vendor specializing in generative AI for engineering design has proposed a cloud-based solution that converts textual descriptions into 3D blueprints. It seems to be a good fit. It is from a well-known Canada-based company with extensive experience providing AI solutions for advanced manufacturing. Billing is based on specific usage, with a fixed amount for each text-to-blueprint generation. It is easy to use, regardless of the technical knowledge of the employees.
- B. **Computer vision tool from a vendor mainly specialized on healthcare:** A vendor with a strong background in AI-driven medical imaging has proposed repurposing its computer vision technology for industrial defect detection. It is a Dallas-based company. The US government is offering some funding to support AI adoption from local organizations. Billing is based on a flat monthly fee, with no limit on usage. There

may be a reduced one-shot fee for an initial proof-of-concept. It is relatively simple to use, but it requires some expert employees to label manufacturing images.

- C. **General cloud tool with AI capabilities:** A large cloud service provider has pitched its AI-powered analytics and automation suite, which includes machine learning models, predictive analytics, and NLP tools. While not tailored specifically for manufacturing, the vendor emphasizes its ability to integrate with ACME AM's existing ERP and IoT systems. It is a European vendor with data centers in Ontario and Alberta. It is a highly technical tool, but the vendor offers free training and certification to your employees, with no limit, including agnostic AI fundamentals courses.

Use the following templates to structure your thought process and prepare an answer to this exercise. The templates will support your exercise, but you can add more information and adapt the sections depending on your own criteria.

**Use case and tool mapping:** This template helps you map the potential use cases and capabilities, and the available tools in a competitive context. That means, if you are evaluating similar tools from different vendors, purely focus on the product features and applications.

| Use case<br>case<br>(a-i) | Use case name                            | Mapping to specific<br>challenges<br>(1-4)                      | Potential tool<br>based on existing<br>information<br>(A-C) | Use case<br>Priority<br>(H/M/L) |
|---------------------------|------------------------------------------|-----------------------------------------------------------------|-------------------------------------------------------------|---------------------------------|
| a                         | AI-powered predictive maintenance system | e.g.,<br><br>1: Process optimization and predictive maintenance | e.g.,<br><br>C: General Cloud Tool with AI capabilities     | (H)igh                          |
| b                         | AI-driven energy optimization            |                                                                 |                                                             |                                 |
| c                         | AI-based demand forecasting              |                                                                 |                                                             |                                 |
| d                         | AI-powered supplier risk assessment      |                                                                 |                                                             |                                 |

|                                                                        |                                               |  |  |  |
|------------------------------------------------------------------------|-----------------------------------------------|--|--|--|
| e                                                                      | Computer Vision-based defect detection        |  |  |  |
| f                                                                      | AI-enhanced quality predictive analytics      |  |  |  |
| g                                                                      | Generative design for product optimization    |  |  |  |
| h                                                                      | AI-driven mass customization platform         |  |  |  |
| i                                                                      | AI-powered digital twin for rapid prototyping |  |  |  |
| <b>Preliminary tool recommendation</b> (before final vendor checklist) |                                               |  |  |  |

By the end of this exercise, you may have a preliminary tool recommendation, but you will need to analyze the vendor information to make sure that the organization is a proper technology partner for your AI initiative. Here's the vendor checklist:

| Category       | Item                      | Comments | Ok/<br>Nok                          |
|----------------|---------------------------|----------|-------------------------------------|
| Vendor profile | Reputation                |          | <input checked="" type="checkbox"/> |
|                | Pace of innovation        |          | <input checked="" type="checkbox"/> |
|                | Geographic considerations |          | ...                                 |
| Tool scope     | Environment               |          |                                     |
|                | Features                  |          |                                     |
|                | Type of models            |          |                                     |
|                | AI knowledge scope        |          |                                     |

|                                         |                         |  |  |
|-----------------------------------------|-------------------------|--|--|
|                                         | Compliance and security |  |  |
| Usage considerations                    | Technical level         |  |  |
|                                         | Interoperability        |  |  |
|                                         | Customization           |  |  |
|                                         | Service level           |  |  |
|                                         | Pricing model           |  |  |
| Key questions                           |                         |  |  |
| Internal considerations                 |                         |  |  |
| Contractual aspects                     |                         |  |  |
| Cost vs value tradeoff                  |                         |  |  |
| Potential ROI                           |                         |  |  |
| Alignment with internal AI usage policy |                         |  |  |
| Final recommendation                    |                         |  |  |
| Name / Signature                        |                         |  |  |

**Worked example:** To show how these templates come together, let's partially complete them for one use case: the AI-powered predictive maintenance system (use case a). In the use case and tool mapping template, it maps to challenge 1 (process optimization and predictive maintenance), and its priority is high, since unplanned downtime directly affects delivery commitments and machinery costs. Looking at the three RFI responses, vendor (A) can be discarded for this use case because text-to-blueprint generation addresses design, not maintenance. Vendor (B) focuses on computer vision for defect detection, which could complement quality use cases but does not cover sensor-based failure prediction. Vendor (C), the general cloud tool, offers the predictive analytics and IoT integration this use case needs, so it becomes the preliminary tool recommendation.

Moving to the vendor checklist for vendor (C), sample observations could look like this: reputation is ok (an established cloud provider with public manufacturing case studies); pace of innovation is ok (frequent releases and a broad partner ecosystem); geographic considerations

require attention (a European vendor with data centers in Ontario and Alberta, so ACME AM must confirm data residency and export control implications for its US aerospace clients); environment and features are ok (a cloud-based suite that integrates with the existing ERP and IoT systems); technical level is not ok at first sight (a highly technical tool), but the unlimited free training and certification offer mitigates the skills gap; and the pricing model requires a proper usage estimation exercise before contract signature.

A justified final recommendation could then read as follows: "We recommend starting with vendor (C) for a predictive maintenance proof of concept, given its fit with challenge 1, its integration capabilities, and the training support that compensates for our current skills gap. Geographic and pricing questions must be resolved during contract negotiation." Your own completed tables will look different, but the logic (map the use case, compare candidate tools, verify the vendor, justify the choice, and sign it) should always be visible.

Remember, this section includes an AI tool checklist for individual analysis, the use case and tool mapping for prioritization, and the final vendor checklist to guarantee you are not getting into a negative client-vendor dynamic.

Rather than treating these tables as forms to fill once, adapt them to your own context: adjust the evaluation factors to your industry and regulatory environment, involve the stakeholders who will live with the decision (procurement, legal, security, finance, and the technical teams), and keep the templates lightweight enough that people actually use them. In larger organizations, you may maintain one version for strategic, company-wide platforms and a simpler one for team-level tools; smaller companies can merge everything into a single one-page document. A few practical tips based on real evaluation processes: designate a clear decision owner for each tool decision, ask them to justify the choice in writing and sign the document, run the checklist with at least two vendors to keep the comparison honest, and revisit completed checklists periodically as vendors and tools evolve. This is especially useful in the context of AI management systems (revisit the end of *Chapter 1* if you don't remember that concept), as companies need to standardize and document the way they interact with external third parties to mitigate risks and to prepare for internal and external audits.

## Leveraging cloud technologies for AI strategy

As we discussed earlier, the multiple pillars of your AI strategy are not isolated streams. Choosing the good people will help you implement and increase the internal AI experience, same as choosing the good tools will accelerate your technical initiatives.

Here is my personal and certainly biased statement: 95% of the companies out there need to rely on public cloud technologies if they really want to make significant progress and accelerate their AI initiative.

If you are not a top technology company (e.g., hyperscaler, AI model developer, very advanced digital native company), then you are part of the AI adopter cohort. And that is okay. You don't need to reinvent the wheel in terms of technology stack, or even AI model development.

This book is not about justifying if you should choose **Amazon Web Services (AWS)**, **Google Cloud Platform (GCP)**, or Microsoft Azure as your main technology partner. Or if you will do better by choosing OpenAI or Anthropic's models. You can evaluate those options and your own company's preferences with your team. But one thing is certain: I haven't seen big organizations, small startups, and even public administrations that make significant progress by building their own technology stack layers. Or at least it is not the case for most of them, and those who succeed were way more advanced than the average company.

Obviously, this will also have an impact on your financial models, as they will be mainly **operational expenditure (OPEX)**-oriented. In practice, this means that AI spending shifts from large upfront capital investments to recurring, usage-based costs that grow with adoption. This changes how you budget and forecast: costs hit your operating budget month after month, they can fluctuate with usage, and they require continuous monitoring and optimization rather than a one-time approval. But you will also reduce the total cost of ownership and allow your teams to focus on value creation rather than complex and long infrastructure deployment projects.

Here are some ideas that justify the impact of cloud adoption on your AI strategy streams, summarized in *Figure 8.2*:




Figure 8.2 – Impact of cloud technologies on AI strategy

Figure 8.2 summarizes the main areas where cloud adoption influences your AI strategy, from financial planning and access to state-of-the-art models to security, governance, and compliance. Let's review each of these impact areas in more detail:

- **Finance:** Financial planning in the cloud allows collaboration between technical stakeholders and finance departments. As part of the cloud economics exercise, the cloud brings new opportunities to balance both **operational expenditure (OPEX)** and **capital expenditure (CAPEX)**. This is part of your company's financial operations (also known as FinOps).
- **ROI friendly:** The value of the cloud for regular IT and data storage activities is clear, but the scenarios where return-on-investment is very clear are those with complex technologies and services that would require heavy upfront investments. For example, advanced research, scientific discovery, and, of course, the newest generative AI technologies. The time and monetary cost of every company building its own platforms would not be sustainable or achievable for most of the organizations.
- **Anti-depreciation:** Access to the latest technologies is expensive, and the underlying hardware components are usually scarce (e.g., GPUs) and lose value fairly quickly. The cloud gives you access to those advanced resources without making heavy investments, especially in physical infrastructure that becomes obsolete very quickly due to the high pace of AI innovation and continuously changing hardware requirements.
- **State-of-the-art AI models:** Cloud migration enables quick and seamless access to the newest **state-of-the-art AI (SOTA)** models, as both serverless APIs (also known as **Model-as-a-Service** or **MaaS**) or/and deploying in the cloud via virtual machines (managed compute). Compared to non-cloud scenarios, this process enables access to AI resources in a simplified manner, with very little deployment or configuration work. This is a clear advantage for companies that cannot afford to hire the science and engineering talent required to deploy AI at scale.
- **Cloud service models:** Scalability and performance are key differentiators of cloud-based AI when compared to on-prem implementations, given the required hardware resources to train and deploy LLMs at scale. Also, the public cloud offers a good balance of control and benefits, thanks to the variety of **IaaS**, **PaaS**, and **SaaS** options (introduced earlier in this chapter, in the *Defining the AI tooling selection criteria* section). The choice and combination of these service models depend on your cloud-enabled AI architectures' required services and functionalities. As a rule of thumb, IaaS gives you maximum control over the environment (e.g., a regulated bank fine-tuning open-source models on dedicated GPU virtual machines) at the price of maximum operational responsibility; PaaS options such as managed ML platforms let your teams build custom

applications without managing servers, trading some flexibility for speed; and SaaS or MaaS options (e.g., ready-made AI assistants, or frontier models consumed via APIs) offer the fastest adoption and lowest operational burden, but the least control and customization. In short, the more managed the service, the lower your operational effort and the higher your dependence on the vendor's roadmap and pricing.

- **Safety and security:** Getting access to the cloud means leveraging native capabilities at different levels: perimeter security, private links to storage accounts, integrated guardrails and countermeasures for AI models, filters for negative or specific topics, etc. These capabilities are not exclusive to AI implementations, but safety and security are certainly key considerations when building AI-enabled applications, and are even more important for generative AI applications.
- **Use case experimentation:** Adopting the cloud includes the advantages of dedicated AI platforms (i.e., workspaces for experimentation and collaboration with low deployment and maintenance burden). This means that the teams can focus on testing and customizing their AI models and architectures to prove value in less time. It also facilitates reusing resources from past use cases and sharing related documentation.
- **Project metrics:** Similar to the experimentation advantage, reducing the platform complexity via cloud allows you to focus on the expected value of your project. You can decide the technical metrics for both development and production phases, and define the "good enough" threshold for your AI implementations. Platforms like Microsoft Foundry (formerly Azure AI Foundry / Azure AI Studio) include pre-defined metrics depending on the type of AI model, e.g., **Receiver Operating Characteristic (ROC)** curve and confusion matrix for classification, RMSE/R-squared for regression, and groundedness for generative AI.
- **MLOps/LLMOps:** If one of the goals of the AI readiness framework is to develop a consistent strategy and evolve the AI experience within the organization, one of the baseline aspects is related to the repeatable, standardized operations related to your AI projects. Disciplines like **Machine Learning Operations (MLOps)** or **Large Language Model Operations (LLMOps)** simplify the approach for the design, development, testing, productization, and maintenance of your AI projects. Cloud platforms include specific features to streamline these processes, which is one of the key differentiators and enablers for AI readiness compared to non-cloud scenarios.
- **Data-driven:** Because of its centralized access and management, cloud platforms help democratize access to data within the organization. Of course, that access is based on roles and policies, previously determined, that will also lead to different roles and responsibilities to own the sources, improve their quality and related metadata, and share the data within the organization.



- **Enablement:** Technology innovation goes beyond core IT and development areas of your company. Regardless of the professional background of the employees, they can access both technical and no/low-code AI tools, enabling everyone to create new applications, test technologies, and understand their capabilities.
- **Accountability:** With that technology enablement and access to new capabilities comes the notion of team and individual accountability. In the cloud, everyone can access information about the cost and impact of their developments.
- **Productivity:** As part of the equation, another key advance of the cloud-enabled tools is access to AI capabilities for day-to-day productivity and technical levels, regardless of the level of professional experience and technology proficiency.
- **Internal data and AI governance:** Getting access to data and AI governance resources based on cloud-enabled tools. Cloud migration enables the organizational foundation, team approach, and Responsible AI practices by providing a set of tools and best practices that support internal governance activities.
- **Compliance:** Same as security and safety measures, compliance is another key advantage of the cloud. The idea is to enable an accountable and non-risk culture that facilitates innovation and is always aligned with data residency, data protection, and new AI laws and standards. Compared to non-cloud scenarios, your company would need to build and document everything themselves.

These are some of the main advantages of cloud adoption in the context of AI. If your company is already leveraging some of these technologies, you can check what's working with your teams as part of the re-evaluation process you will need during the annual or multi-year negotiations with hyperscalers. You can leverage the things that are not working either as potential recommendations for improvements (e.g., increasing support speed when something goes wrong) and as levers for better negotiations (e.g., obtaining discounts if your provider doesn't reach specific KPIs or SLAs).

## Building resilience while avoiding vendor lock-in

The AI ecosystem is evolving at a pace that makes single-vendor strategies increasingly fragile. Capabilities leapfrog one another across providers, pricing models shift, regulatory environments change and harden, and geographic constraints multiply. What is state-of-the-art today may be economically irrational tomorrow. In this environment, committing entirely to one provider is rarely a healthy long-term strategy.

Over time, most organizations find themselves working with multiple providers. They do so for sensible reasons. Cost structures vary dramatically depending on workload type, latency requirements, and scale. One provider may excel in reasoning, another in multimodal

performance, another in domain-specific fine-tuning. Also, risk management considerations make concentration dangerous (e.g., regulatory and data residency requirements may demand new regional deployments that are not available in all providers). What begins as tactical diversification gradually becomes structural complexity.

The strategic mistake is framing the problem as an AI model comparison exercise. That's only a short-term factor in a long-term game because models improve and benchmarks change. The more durable questions are architectural: Can we avoid hard lock-in? Can we evaluate providers consistently? Can we switch models without breaking products or destroying economics? Can we govern coherently across heterogeneous systems? The goal here is to have some degree of flexibility.

Achieving this requires a deliberate architectural stance. Applications should not be tightly coupled to a specific model version, a proprietary API schema, or a vendor-specific prompt structure. When product logic embeds provider assumptions deeply into the codebase, every vendor change becomes a costly rewrite. By contrast, introducing an abstraction layer separates application logic from model and platform selection. This architectural decision transforms switching providers from a structural change into a configuration decision.

## Technical flexibility and governance

In the technical aspect, a practical implementation of this is the idea of a model gateway. Rather than allowing product and engineering teams to integrate directly with multiple providers, requests flow through a centralized control plane. This layer standardizes inputs and outputs, manages credentials, tracks usage and cost, applies safety filters, and routes traffic based on policy. Routing decisions may consider task type, latency sensitivity, regional requirements, regulatory constraints, availability, or real-time cost optimization. The gateway becomes both an operational safeguard and a strategic lever.

However, technical flexibility alone is insufficient. Multi-provider environments introduce governance complexity that cannot be outsourced. Providers differ in data retention policies, training data usage terms, safety filtering mechanisms, and service-level guarantees. If governance standards vary implicitly by vendor, risk becomes invisible and unevenly distributed. Organizations must define internal safety, compliance, and documentation standards that sit above the provider layer. Evaluation frameworks must be standardized and provider-agnostic, observability and logging should aggregate centrally, and vendors should be treated as strategic third-party risk entities.

Also, economic discipline becomes especially critical in multi-provider environments. Without centralized visibility, cost drift accelerates. Teams default to the highest-performing models regardless of marginal value. Token consumption expands without constraint. Experimentation

multiplies across silos. Over time, AI spending becomes unpredictable and difficult to attribute to business outcomes.

For this purpose, the idea of FinOps for AI is especially relevant. Preventing this requires explicit financial operations design. Major workflows should have defined cost-per-transaction targets. Lower-complexity tasks should be routed to lower-cost models when performance thresholds are met. Model selection should be informed by measurable performance-to-cost trade-offs, not anecdotal preference. Budget ownership and usage transparency must be built into the control plane itself.

There is also an organizational dimension. Multi-provider strategy introduces coordination challenges. Without clear ownership, model sprawl emerges. Different teams build parallel abstractions. Redundant evaluations proliferate. Vendor relationships become fragmented. To prevent this, model governance must be institutionalized. A centralized AI platform function should define standards, maintain evaluation benchmarks, oversee routing policies, and negotiate vendor contracts. This function does not slow innovation; it prevents entropy.

In practice, institutionalizing this function usually means creating a small central AI platform team, or extending an existing platform engineering or center of excellence structure, a pattern that many large enterprises have publicly described and that analyst firms regularly document in their research on AI operating models. A core team of a few platform engineers and architects owns the model gateway, publishes the approved model catalog, maintains shared evaluation benchmarks, and manages the vendor relationship calendar (renewals, roadmap reviews, escalations). The key is to give this team an explicit mandate and executive sponsorship: product teams keep their freedom to build, but they consume models and providers through the shared platform rather than signing their own contracts. Starting small works well; even one architect and one engineer maintaining a gateway, a benchmark suite, and a contract inventory can prevent most of the sprawl described earlier.

## **Negotiation, exit strategies, and open-source alternatives**

Vendor negotiation deserves the same level of discipline. When negotiating AI contracts, a few tactics consistently pay off: negotiate before you are dependent, while you still have credible alternatives; ask for price protection clauses and caps on unilateral pricing changes, since AI pricing models are still volatile; secure commitments on data portability, including the right to export your data, fine-tuned models (or at least the training data behind them), and usage logs in standard formats; align contract duration with the pace of the market, favoring shorter terms or renegotiation windows over long lock-ups; and tie part of the payments to service-level agreements and support quality, as mentioned earlier when discussing negotiation levers with hyperscalers.

What if you are already locked in? The exit playbook starts with an honest inventory: identify where provider-specific assumptions live in your systems (API schemas, prompt structures, proprietary features, and data formats) and estimate the switching cost for each. Then reduce that cost progressively rather than attempting a big-bang migration: introduce the abstraction layer or gateway in front of the incumbent provider first, refactor the most critical applications to use it, and run a second provider in parallel for a subset of low-risk workloads to validate performance and build internal confidence. Your negotiation leverage improves as your credible ability to leave increases, so even a partial migration path changes the dynamics of the next renewal conversation.

Open-source and self-hosted alternatives are another risk mitigation lever in this equation. Open-weight models that you can deploy on your own infrastructure (or on rented cloud GPUs), as well as local inference options for smaller models, give you a fallback that no commercial contract can remove. Not every organization will run these options in production, and they carry their own operational costs, but keeping an evaluated open-source alternative in your model catalog acts as an insurance policy: it caps the price a vendor can credibly charge you, and it guarantees continuity if a provider changes its terms, deprecates a model, or exits your region.

Designing for multiple providers is not about hedging indecision. It is about preserving strategic leverage in a rapidly shifting ecosystem. Optionality increases negotiating power. It protects against outages and pricing shocks. It accelerates adoption of new capabilities. But doing this without structure leads to technology chaos in the form of fragmented architectures, uncontrolled costs, inconsistent governance, and operational fragility.

Remember, the organizations that will lead in AI will not be those that chose one single model or vendor. They will be those that designed systems capable of evolving in the next few years without adding a lot of technical burden. In this case, flexibility is a strategic asset, but some degree of discipline is what will make your AI initiative sustainable.

## Summary

In this chapter, you learned how to navigate the AI ecosystem from a leader's perspective. We started by demystifying the AI technology stack, exploring how infrastructure, platforms, development studios, libraries and frameworks, and governance layers build on each other to turn raw compute into working AI products.

We then moved to the evaluation of tools and vendors, where you reviewed a practical set of selection criteria and applied them through the ACME Advanced Manufacturing exercise, using the individual tool checklist, the use case and tool mapping, and the vendor checklist templates. We also discussed why most organizations should leverage cloud technologies to accelerate their AI strategy, and what that means for your financial models and total cost of ownership. Finally, we covered how to build resilience while avoiding vendor lock-in, including abstraction layers and model gateways, FinOps discipline, contract negotiation tactics, and open-source alternatives. These skills will help you turn technology choices into strategic leverage rather than long-term constraints.

In the next chapter, we will focus on AI governance, compliance, and risk management, so you can keep this growing technology landscape trustworthy and under control.

# 9

## Responsible AI Implementation: Ethics, Governance, and Accountability

From an AI leadership perspective, advancing your initiative and products in a responsible and tangible way is not optional. It is not even about being ethical and doing things right because of your individual and organizational willingness to do so. It is something that goes beyond the opportunity to do things right, and enters a dimension of corporate responsibility, reputation, and actual accountability through regulations and fines.

Often perceived as an operational and financial burden, the goal of this chapter is to demystify what ethical approaches and being responsible mean in the context of AI. In general terms, you can do things right up front and proactively, or wait until it is too late and you have to do it because regulations, industry (you and your competitors), or even end users say so.

In this final chapter, you will learn how to move from abstract ethical intentions to a working Responsible AI practice: distinguishing ethical, responsible, accountable, and trustworthy AI, quantifying your initiatives with concrete metrics and tools, navigating the growing landscape of AI regulations and standards, and implementing data and AI governance in your organization. These learnings matter because they will help you protect your company from legal, financial, and reputational risks while building the trust that customers, employees, and regulators increasingly demand. By the end of the chapter, you will be able to design a governance approach that turns responsibility from a perceived burden into a real competitive advantage.

We will cover the following key topics:

- From AI ethics to accountability

- Quantifying Responsible AI initiatives
- Relation between ethics and regulations
- Notion of data and AI governance

Let's start with the fundamentals!

## From AI ethics to accountability

The concept of *AI ethics* is important, and its significance to social well-being, human rights, and the societal impact of AI is undeniable nowadays. With a wider adoption of AI tools by the general public, ensuring an ethical use of these technologies is key to avoiding their potential negative outcomes.

However, there is no single or universally accepted definition of "ethics" in general, or what is considered "ethical," as these interpretations are shaped by various factors, including moral values related to specific geographies or an organization's cultural norms, traditions, beliefs, and religions. For example, the definition of ethical AI for Europe or Canada is probably not the same as it is for companies like Twitter/X, TikTok, or Palantir. Each of them have their specific contexts, overall leadership and strategy styles, and business activities.

Unlike the regulatory frameworks, which are legally binding, ethical guidelines are not enforced by government authorities or compliance requirements. Instead, they consist of voluntary principles that individuals and organizations can choose to adopt, especially if they resonate with their own values.

And, of course, the debates surrounding data and AI are not just about the technical capabilities, but also about its potential to cause harm or benefit. While regulations address legal issues such as data privacy and security, they do not fully encompass the ethical dimension. Moreover, because ethics are not universally defined, determining whether a data application is ethical is not always straightforward.

That said, both AI ethics and regulations are currently converging to bring something that would be both legal and desirable, as the optimal scenario for any AI implementation.

AI governance can be implemented in a variety of ways, ranging from broad intentions to concrete, actionable steps. Concretely:

- **Ethical AI:** This level is concerned with embedding core moral values, such as fairness, transparency, privacy, and respect for human rights, into the design and operation of AI systems. It aims to ensure that AI technologies support societal values and avoid causing harm or reinforcing bias. Organizations around the world, including governments, companies, and research institutions, develop their own sets of ethical principles to guide AI development. For inspiration, explore various principle frameworks compiled

by the Principled AI initiative at Harvard University's Berkman Klein Center and consider developing a customized set of ethics for your organization.

- **Responsible AI:** Building on the foundation of ethical AI, **Responsible AI (RAI)** involves actively managing and governing AI systems throughout their lifecycle, from design and development to deployment and monitoring. This approach focuses on applying best practices and technical solutions to ensure AI is safe, inclusive, and aligned with human well-being. It represents a more operational step toward embedding ethical intentions into real-world systems.
- **Accountable AI:** At this stage, the focus shifts to establishing clear oversight, ownership, and traceability for AI systems. The goal is to ensure that individuals or organizations can be held responsible for how AI is developed and used, and for its impact on people. This includes adhering to relevant legal and regulatory standards and implementing robust mechanisms for transparency and accountability.
- **Trustworthy AI:** This is the overarching objective that unites the previous three approaches. Trustworthy AI refers to systems that consistently demonstrate reliability, fairness, transparency, and alignment with human values, earning and maintaining public trust. Ethical, Responsible, and Accountable AI are all integral to achieving this goal.

These terms represent incremental levels of action that any individual, company, or organization can implement depending on their moral values, resources, and internal AI governance and maturity.

## Key terms and principles

As the global conversation around the ethics of artificial intelligence deepens, various frameworks have emerged to guide its responsible development. Among the most influential are the Asilomar AI Principles, created by leading researchers and practitioners to ensure AI is aligned with human values, and the Principled AI collection curated by Harvard University's Berkman Klein Center (<https://cyber.harvard.edu/publication/2020/principled-ai>), which compiles diverse ethical guidelines from around the world. That collection contains several initiatives from countries, cities, companies, and supranational organizations.

Another example is the Montreal Declaration for a Responsible Development of Artificial Intelligence (<https://montrealdeclaration-responsibleai.com/the-declaration/>). It is a foundational document that outlines ethical principles to guide the development and deployment of AI technologies, based on the very own "personality" of the Montreal society.

Initiated by the Université de Montréal, it was the result of an inclusive deliberation process involving citizens, experts, public officials, industry stakeholders, civil organizations, and professional associations. The Declaration aims to ensure that AI serves the common good and



respects fundamental human rights. Concretely, it is structured around ten ethical principles, each reflecting a commitment to aligning AI development with societal values.

Rather than reproducing the official description of each principle, which you can read in full at the link provided above, here is a condensed view of the ten principles and what each one implies for your AI initiatives:

1. **Well-being:** AI should support human flourishing, promoting mental, emotional, social, and physical health rather than exploiting behavioral vulnerabilities.
2. **Respect for autonomy:** AI must enhance individual agency, keeping users informed, in control, and able to opt out of AI-based decisions.
3. **Protection of privacy and intimacy:** AI systems must minimize data collection, apply privacy-preserving techniques, and refrain from invasive surveillance practices.
4. **Solidarity:** AI should foster community well-being and distribute risks and benefits fairly across society, including marginalized groups.
5. **Democratic participation:** AI development and deployment should remain open to public scrutiny, input, and challenge, especially for high-risk applications.
6. **Equity:** AI must actively counteract structural inequalities instead of reproducing or amplifying them.
7. **Inclusion and diversity:** Development teams, datasets, and user experiences must represent the full spectrum of human identity and experience.
8. **Prudence:** Decision-makers must anticipate downstream harms and prefer cautious, reversible decisions in high-stakes or high-uncertainty areas.
9. **Responsibility:** There must be clear lines of legal, ethical, and operational accountability for the actions of AI systems.
10. **Sustainable development:** AI should minimize its environmental footprint and contribute to the long-term health of the planet.

Regardless of the nature of the initiative, all of them focus on key principles related to proper respect for human rights and the fair expectation of a positive impact for users and society in general. Some additional terms you need to know:

- **Transparency:** This is one of the most important principles, and it refers to the openness of an AI system's design and function, ensuring that all stakeholders can access meaningful information about how the system was built, what data it uses, how it makes decisions, and what its limitations are. This enables scrutiny, fosters public trust, and is a foundation for ethical oversight.

- **Explainability:** It focuses on the ability of an AI system to generate understandable reasons for its decisions or behavior. This means users can know why the AI made a specific decision and receive a response they can make sense of.
- **Interpretability:** This is different from explainability, and it involves the technical clarity of a model's inner workings so that developers and experts can examine, troubleshoot, and refine the system. For instance, simpler models like decision trees are inherently interpretable, while complex ones like neural networks require specialized techniques to unpack their behavior.
- **Bias reduction:** This entails actively identifying, mitigating, and monitoring biases in the training data, algorithm design, and deployment context of AI systems. The goal is to prevent unfair outcomes such as discrimination against marginalized groups, which can arise from historical data imbalances or flawed assumptions. This is a continuous process requiring fairness audits, diverse teams, and inclusive testing.
- **Accountability:** Related to the accountable AI term, this establishes clear lines of responsibility for AI outcomes, meaning that individuals or organizations can be held liable for decisions made by or with AI systems. This involves creating audit trails, governance structures, and escalation procedures so harms can be investigated and addressed, especially in the case of errors or misuse.
- **Inclusiveness and fairness:** It refers to the design and deployment of AI systems that respect diversity and promote equal access and outcomes across different demographics. This means engaging with a wide range of users during development, ensuring accessibility, and testing for disparate impact to avoid reinforcing existing inequalities.

These principles are not theoretical, and some publicly reported cases show what happens when they are overlooked. Amazon reportedly discontinued an experimental AI recruiting tool after discovering that it systematically downgraded resumes associated with female candidates, a direct illustration of why bias reduction and inclusive training data matter. Similarly, the COMPAS recidivism scoring system used by some US courts became the center of a public debate after investigative journalists reported that its error rates appeared to differ across racial groups, highlighting the importance of transparency and independent fairness audits for high-stakes decisions.

On the explainability side, the Apple Card attracted a regulatory investigation in 2019 after customers publicly complained about apparently inconsistent credit limits between spouses. Regardless of the final outcome of that investigation, the reputational lesson was clear: if your organization cannot explain how an algorithmic decision was made, your customers and regulators will assume the worst. Each principle in the list above maps to very tangible business, legal, and reputational risks.

Let's now see how we can make the ethical approach more tangible by choosing our areas of action for principles and then exploring related tools and metrics.

## Quantifying RAI initiatives

Ethical AI principles usually sound great to any organization or individual, but they are meaningless if they are not actually implemented in the day-to-day of the company. The only way to make sure that's the case is by quantifying specific results that can also support internal or external audit activities. Here is an illustrative selection of metrics and tools that could enable you to implement measurable actions:

| Principle      | Approach                   | Description                                                  | Example metrics                                      | Technical tools                                                            |
|----------------|----------------------------|--------------------------------------------------------------|------------------------------------------------------|----------------------------------------------------------------------------|
| Transparency   | System documentation       | Document models, data, and decision processes for visibility | Model card completeness, documentation coverage      | Model cards, <b>Weights and Biases (W&amp;B)</b>                           |
|                | Auditability               | Enable reviewability and traceability of decisions           | Time to trace decision, audit readiness index        | Fiddler AI, WhyLabs, Snowflake AI Observability (formerly TruEra)/ TruLens |
| Explainability | User comprehension testing | Evaluate if users understand the system's outputs            | Survey scores, user accuracy in predicting decisions | Zeno, general UX testing tools, custom usability studies                   |
|                | Explanation fidelity       | Align explanations with true model behavior                  | SHAP, LIME scores                                    | SHAP, LIME, Captum                                                         |

|                      |                               |                                                                |                                             |                                                |
|----------------------|-------------------------------|----------------------------------------------------------------|---------------------------------------------|------------------------------------------------|
| Interpretability     | Model simplicity scoring      | Score model complexity and interpretability                    | # of parameters, tree depth                 | Skater, InterpretML                            |
|                      | Simulatability                | Measure whether humans can understand and anticipate AI output | Human-AI prediction alignment               | Custom evaluation tools, A/B testing platforms |
| Bias reduction       | Fairness audits               | Detect and mitigate disparities across demographic groups      | Disparate impact, equalized odds            | AI Fairness 360, Fairlearn, Z-Inspection       |
|                      | Subgroup performance analysis | Evaluate how performance varies across sensitive groups        | False positive/negative gaps                | What-If Tool, Microsoft Responsible AI Toolbox |
| Privacy preservation | Differential privacy modeling | Apply mathematical guarantees to protect personal data         | $\epsilon$ (epsilon) in DP, privacy leakage | OpenDP, PySyft, SmartNoise                     |
|                      | Data access controls          | Monitor who can see and use personal data                      | # of unauthorized access attempts           | Collibra, Immuta                               |
| Accountability       | Governance frameworks         | Assign responsibility across the AI lifecycle                  | RACI coverage, policy compliance score      | Monitaur, Credo AI, OneTrust                   |

|                            |                                    |                                                     |                                                          |                                                                                        |
|----------------------------|------------------------------------|-----------------------------------------------------|----------------------------------------------------------|----------------------------------------------------------------------------------------|
|                            | Incident resolution tracking       | Detect and manage harms or adverse events           | # of AI-related incidents logged and resolved            | Custom incident tracking tools                                                         |
| Inclusiveness and fairness | Inclusive development processes    | Ensure diverse perspectives are involved in design  | Diversity index, participatory session frequency         | Diverse Teams Checklists, inclusive design practices, human-centered design frameworks |
|                            | Accessibility and usability checks | Make systems usable by all populations              | WCAG compliance, assistive tech compatibility            | Deque Axe, WAVE, Fable Tech Labs                                                       |
| Robustness and safety      | Stress and failure testing         | Evaluate AI under extreme or adversarial conditions | Adversarial test success rate, error under perturbation  | Robustness Gym, CleverHans                                                             |
|                            | Redundancy and fallback systems    | Ensure graceful degradation or failover options     | System recovery time, uptime %                           | Chaos Engineering tools, container health checks                                       |
| Human oversight            | Human-in-the-loop design           | Ensure humans can review and override AI decisions  | % of override opportunities, manual intervention latency | Custom review platforms, approval pipelines                                            |

|                      |                               |                                                            |                                                        |                                                |
|----------------------|-------------------------------|------------------------------------------------------------|--------------------------------------------------------|------------------------------------------------|
|                      | Escalation protocols          | Provide clear pathways for human intervention              | Escalation success rate, clarity index                 | ServiceNow, Jira, custom escalation procedures |
| Sustainability       | Energy and resource tracking  | Monitor environmental footprint of AI training/deployment  | Carbon emissions, GPU/TPU hours, energy per inference  | CodeCarbon, Carbontracker                      |
|                      | Life cycle impact assessments | Assess sustainability throughout AI system lifecycle       | Resource consumption efficiency score                  | Green Algorithms, LCA tools                    |
| Well-being           | Impact assessments on users   | Evaluate psychological, social, and health effects         | Stress level reports, user well-being score            | Surveys, HCI-based UX platforms                |
|                      | Harm minimization checks      | Proactively identify and reduce potential negative impacts | # of avoided harms, risk categorization counts         | Z-Inspection, internal risk rating tools       |
| Autonomy and consent | Consent mechanisms            | Ensure users provide informed and revocable consent        | % of informed opt-ins, readability of consent language | Custom UI modules                              |

|  |                                        |                                                                              |                                                                           |                                                                        |
|--|----------------------------------------|------------------------------------------------------------------------------|---------------------------------------------------------------------------|------------------------------------------------------------------------|
|  | Decision-<br>support vs.<br>automation | Balance autonomy<br>by clarifying when<br>AI is assistive or<br>substitutive | % of<br>automated vs.<br>assistive<br>decisions,<br>override<br>frequency | UX toggles, AI<br>Disclosure<br>Tools,<br>Explainable UI<br>components |
|--|----------------------------------------|------------------------------------------------------------------------------|---------------------------------------------------------------------------|------------------------------------------------------------------------|

Table 9.1: An illustrative selection of metrics and tools to implement measurable actions

Of course, as an AI leader within your organization, you are not expected to know all these tools and metrics. Actually, not even your technical team members will know all this. But the table can serve as an illustrative example for you and your company to find specific ways to prove and quantify the progress of your AI ethics, depending on the principles and main areas of action. Feel free to add other principles, approaches, and tools to your responsible AI mix.

## Relation between ethics and regulations

The relationship between AI ethics and AI regulation reflects a crucial duality in the governance of generative AI: the first one is voluntary, grounded in moral values and professional responsibility; the other is mandatory, backed by legal authority and formal enforcement. These two dimensions are not in conflict, as they complement and reinforce each other in shaping how AI is developed, deployed, and governed.

AI ethics consists of normative frameworks that articulate what should be done to ensure AI is beneficial, fair, and aligned with human values. These guidelines are typically published by academic institutions, international organizations, or private companies, and serve to inspire responsible innovation and prevent social harm. However, because they are voluntary, ethical principles often lack accountability or enforcement mechanisms. As a result, there's a risk of ethics-washing, where organizations showcase ethical commitments publicly but fail to implement them meaningfully in practice.

On the other hand, AI regulation involves legally binding rules enforced by governments or regulatory bodies. These rules require organizations to comply with minimum safety, transparency, and fairness standards and are especially crucial in high-risk sectors such as healthcare, law enforcement, or employment. Regulatory frameworks often impose sanctions for non-compliance, such as fines or operational restrictions, which creates a strong incentive for companies to adhere to them.

While ethics tends to be more proactive and adaptable, regulation is often reactive, emerging after public concerns, technological developments, or ethical failures. Nonetheless, the two interact dynamically: ethical discourse frequently informs the creation of regulation, and

regulations can embed ethical standards into enforceable law. Together, they form a holistic approach to AI governance.

## Types of regulations

As an AI leader and adopter, you need to understand the mix of regulations and requirements that will guide your responsible AI approach. The mix will consist of a combination of laws, based on the scope of application, and the specific geography where your company is operating.

The idea is to treat compliance as a multi-layered responsible approach, integrating ethical design, technical controls, and legal adherence upfront. Failure to align with these regulations can result in legal penalties, reputational harm, and loss of public trust. Also, adopting them in a proactive manner is always better than doing so reactively.

**AI-specific regulations:** These are emerging (existing or work in progress, depending on the geography) legal frameworks developed to address the unique risks of AI itself, including opacity, autonomy, and systemic bias. They aim to provide horizontal governance across sectors for AI applications. Some examples include the following:

- **European Union's AI Act:** The EU AI Act is the world's first comprehensive legal framework that categorizes AI systems by risk and imposes strict obligations on high-risk applications.
  - **Approach:** It classifies AI systems by risk levels (unacceptable, high, limited, minimal), with specific focus on high-risk use cases that must meet stricter requirements, especially for areas like healthcare, policing, and education.
  - **Core obligations:** The high-risk systems must include:
    - **Risk management systems:** Organizations must establish structured processes to identify, assess, monitor, and mitigate risks associated with AI systems throughout their lifecycle.
    - **High-quality datasets:** AI systems must be trained and validated using accurate, representative, and relevant data to ensure fairness, reliability, and effectiveness.
    - **Record-keeping and transparency:** Developers and deployers must maintain detailed documentation of system design, data sources, decision logic, and testing procedures to ensure auditability and regulatory compliance.



- Human oversight: AI systems, especially in high-impact contexts, must include mechanisms for meaningful human supervision, intervention, or override to prevent harm and ensure accountability.
- Robust performance and cybersecurity: AI must be engineered to perform reliably under a range of conditions and be protected against adversarial attacks, data breaches, and system manipulation.

Canada currently has no horizontal AI law after AIDA died with Bill C-27 in January 2025, so companies operating in both regions should treat the EU AI Act as the effective baseline.

- **Colorado's AI Act** : Colorado's AI Act (SB 24-205) was the first broad US AI consumer-protection law, but it never took effect: it was repealed and replaced by SB 26-189 (signed May 2026, effective 1 January 2027), which drops the duty of care, risk management program and impact assessments in favor of developer documentation and consumer disclosure.
  - **Approach**: A risk-based model focused on "high-risk AI systems". In general (and not very different from the European approach), this means AI that makes or is a substantial factor in making consequential decisions (e.g., employment, housing, lending, insurance, healthcare). Obligations attach to both developers and deployers (the organizations using the system).
  - **Core obligations (high-risk AI)**:
    - Reasonable care to prevent algorithmic discrimination (developer + deployer duty), with a "safe harbor"-like rebuttable presumption if specified steps are followed.
    - Developer disclosures + documentation support: Provide deployers with key system information and documentation needed for impact assessments, plus public-facing summaries of the types of high-risk systems offered and risk-management practices; notify the AG/deployers of known/foreseeable discrimination risks within set timeframes.
    - Deployer governance + impact assessments: Implement a risk management program, conduct impact assessments, and perform at least annual reviews to check for algorithmic discrimination.
    - Consumer transparency + recourse: Notify consumers when a high-risk AI system is used as a substantial factor in a consequential decision; offer the ability to correct data and appeal adverse decisions, including human review when technically feasible.

- Public transparency + incident reporting: Publish summaries of high-risk AI use and risk management; report discovered algorithmic discrimination to the Colorado AG within defined windows.
- Notable carve-outs/exemptions: Some deployers (e.g., small employers under certain conditions) may be exempt from parts of the risk management/impact assessment/public disclosure obligations.

Colorado initially moved toward the EU model, but SB 26-189 (2026) repealed those obligations in favor of developer documentation and consumer disclosure, so do not assume US-EU convergence; design for the stricter EU baseline and layer state disclosure duties on top.

- **China's AI Regulation Framework:** China's AI governance framework includes strict rules for algorithm transparency, content moderation, and ethics review, especially for recommender and generative AI systems. Similar to the rest of regulations, adopters may face fines, suspension, or market bans for non-compliance.
  - **Approach:** It includes mandatory registration, content moderation, and ethics assessments for AI systems like recommendation algorithms and generative models.
  - **Core obligations:** The implications for AI adopters:
    - Need to implement technical and organizational safeguards, like explainability, risk assessments, and documentation.
    - Must engage in ongoing compliance processes, especially for high-risk AI

Interestingly enough, China was the first to create a series of interim measures for the application of generative AI technologies, before any other geography around the world.

- **Canada's Artificial Intelligence and Data Act (AIDA):** AIDA was a proposed federal framework for regulating artificial intelligence. Introduced as part of Bill C-27, it would have established requirements for certain AI systems, particularly those considered "high-impact." The proposed framework focused on issues such as risk assessment, mitigation of biased outcomes, transparency, record-keeping, and regulatory oversight. However, Bill C-27 did not become law, and AIDA therefore never took effect. Under the proposed AIDA framework, organizations responsible for high-impact AI systems would have been expected to identify and assess risks of harm or biased output, establish measures to mitigate those risks, maintain appropriate records, and meet transparency requirements. The proposal also contemplated regulatory oversight and substantial penalties for serious non-compliance. These requirements should therefore be understood as features of a proposed regulatory regime rather than obligations currently imposed by AIDA.

**What applies in Canada today:** In the absence of AIDA, AI systems in Canada remain subject to existing federal and provincial laws and sector-specific regulatory requirements. Of particular importance are the following:

- **Personal Information Protection and Electronic Documents Act (PIPEDA):** Canada's federal private-sector privacy law governs how organizations subject to the Act collect, use, and disclose personal information in commercial activities. Its principles include accountability, consent, limiting collection and use, safeguards, openness, individual access, and mechanisms for challenging compliance. These requirements can apply when AI systems process personal information.
- **Quebec's Law 25:** Quebec's modernized private-sector privacy regime contains requirements particularly relevant to AI and automated processing of personal information. These include transparency obligations concerning decisions based exclusively on automated processing and privacy impact assessment requirements in specified circumstances. Organizations deploying AI systems involving personal information in Quebec must therefore consider these obligations even though Canada does not currently have a general AIDA-style federal AI law.
- **OSFI Guideline E-23 – Model Risk Management:** For federally regulated financial institutions, the **Office of the Superintendent of Financial Institutions (OSFI)** published the final Guideline E-23 in 2025. The guideline expressly encompasses models using AI and machine-learning methods and establishes expectations for enterprise-wide, risk-based model governance throughout the model lifecycle. Among other things, it addresses model identification, risk ratings, review, deployment, monitoring, documentation, and explainability. The final guideline is scheduled to take effect on May 1, 2027.

Canada's current approach to AI governance is therefore better understood as a combination of existing privacy law, provincial requirements, and sector-specific regulation rather than a single comprehensive federal AI statute. Organizations should determine which requirements apply based on factors such as their jurisdiction, industry, use of personal information, and the purpose and impact of the AI system.

Other geographies are still deciding their approaches, but you can expect new regulatory approaches as AI adoption continues to evolve. If you want to keep track of new AI (or AI-related) regulations, you can use Fairly.ai's global interactive map (<https://www.fairly.ai/blog/map-of-global-ai-regulations>).

**Data protection and privacy regulations:** These are foundational (and already existing) laws that apply to any AI system processing personal data, regardless of industry or use case. They aim

to protect individuals' rights to privacy, transparency, and control over their data. Their scope of application and key principles depend on the regulation, but some common areas are:

- **Informed consent for data collection and use:** Individuals must be clearly informed about what data is being collected, how it will be used (including in AI systems), and must give explicit, voluntary consent before processing begins.
- **Right to access, correct, or delete personal data:** Users have the right to view the personal data held about them, request corrections to inaccurate information, or demand deletion when appropriate (e.g., if consent is withdrawn).
- **Restrictions on automated decision-making:** Laws often limit the use of AI to make decisions with significant impacts (e.g., hiring, credit, legal outcomes) unless safeguards like human review or user rights to contest are in place.
- **Data minimization and purpose limitation:** Only the minimum amount of personal data necessary should be collected, and it must be used solely for the specific, declared purposes, which implies not reusing it without renewed consent.

Some examples of data protection and privacy regulations include these local and international laws:

- **European Union's General Data Protection Regulation (GDPR):** The EU's comprehensive data privacy law grants individuals strong rights over their personal data, including access, correction, erasure, and restrictions on profiling among others, while requiring organizations to ensure transparency, lawful processing, and accountability. GDPR was the baseline for other international regulations that ended converging with the European approach. There is also a GDPR UK version that shares the same core principles.
- **California's Consumer Privacy Act/Privacy Rights Act (CCPA/CPRA):** These laws provide California residents with rights to know, access, delete, and opt out of the sale or sharing of their personal data, with the CPRA adding stricter rules on sensitive data and automated decision-making. In the absence of a federal-level regulation in the US, the Californian approach served as the de facto data privacy standard for American companies, given California's market size and the concentration of large technology firms in the state.
- **Canada's Personal Information Protection and Electronic Documents Act (PIPEDA):** Canada's federal privacy law governs how private-sector organizations collect, use, and disclose personal information, requiring meaningful consent, access rights, and safeguards for data used in AI and other digital processes.

The implications for AI adopters like you include requirements for AI systems that use personal data (e.g., for profiling or recommendation systems like Netflix or Amazon), as they must be

explainable and allow for human review. There is also an important term that you should know: **Data Protection Impact Assessments (DPIAs)**, which are a structured process required under these data protection laws to evaluate the risks to individuals' privacy when an organization plans to process personal data via AI systems, especially for high-risk AI use cases.

To picture how this structured process takes shape in the real world, imagine an illustrative example of a bank that wants to deploy an AI model to pre-score consumer loan applications. The data protection team would first describe the processing: what personal data is collected (income, employment history, transaction patterns), from which sources, and for what purpose. They would then assess necessity and proportionality, asking whether each data point is really needed for the scoring decision. Next, they would identify the risks to individuals, such as unfair exclusion from credit, re-identification of applicants, or excessive data retention, and define mitigation measures such as data minimization, pseudonymization, human review of rejections, and clear retention limits. Finally, the assessment would be documented and signed off by the data protection officer before deployment, and revisited whenever the model or its data sources change. The exercise creates a reusable record that regulators can request at any time.

- **Vertical sector-specific regulations:** These are legal and regulatory requirements tailored to specific industries where the consequences of AI errors, misuse, or bias are particularly critical, such as in healthcare, finance, transportation, and employment. In these sectors, AI systems often interact with sensitive data, life-and-death decisions, or high-impact outcomes, and therefore require stricter safeguards, validations, and oversight.
  - **Healthcare:** In general terms, AI in healthcare must comply with strict data privacy laws and undergo clinical validation to ensure safety and effectiveness in diagnostics and treatment. Some applicable regulations include the US Health Insurance Portability and Accountability Act. HIPAA protects patient data privacy and governs how AI can access or process **electronic health records (EHRs)**. Even if it is local to the United States of America, this regulation is a standard for any health data protection topic around the world. EMA (EU) and Health Canada also issue AI guidance on clinical trials and digital health.
  - **Finance:** Financial AI systems, like credit scoring or fraud detection, must adhere to regulations to ensure fairness, transparency, and auditability. Applicable regulations include Basel III (international), FCRA (US), PSD2 (EU), and SEC/FINRA compliance. The Canadian OSFI finalized Guideline E-23 on Model Risk Management, which explicitly covers AI and machine learning models at federally regulated banks and insurers and takes effect on 1 May 2027.
  - **Transportation and mobility:** AI used in autonomous vehicles must comply with safety standards and government regulations, ensuring reliability, risk

assessment, and ethical decision-making, and guaranteeing the protection of physical safety. Some examples include ISO 26262 for functional safety of automotive electrical/electronic systems, complemented by ISO 21448 (SOTIF) and ISO/PAS 8800:2024, which specifically addresses safety and artificial intelligence in road vehicles.

- **Employment and labor:** AI in hiring and employee monitoring must comply with equal opportunity laws and undergo bias audits to prevent discrimination and promote fairness. For example, Canada's Human Rights Act and Labor Codes apply where AI impacts employment equity, and New York Local Law 144 established bias audits for hiring tools.

The main complexity is to harmonize existing vertical regulations with new AI laws, and to do so in a way that does not contradict any of them. Other implications include:

- **Human oversight:** Many sectors mandate human-in-the-loop supervision, especially for high-stakes decisions (e.g., diagnoses, loan approvals) that are also subject to high-risk implications in regulations such as the EU AI Act.
- **Audit readiness:** This implies the need to maintain documentation, bias impact statements, and technical explainability records for regulators.
- **Adaptability:** AI systems must be designed with flexibility to adjust to evolving regulations, especially in dynamic fields like digital health or finance.
- **Certifications and testing:** Certain sectors may require third-party validation or certifications before AI can be deployed (e.g., FDA clearance or ISO audits).

To make this regulatory mix more tangible, consider how it would play out for three common AI projects:

- **A healthcare diagnostic AI system:** A model that helps radiologists detect anomalies would typically qualify as high-risk under the EU AI Act, requiring risk management, high-quality datasets, human oversight, and detailed documentation. Because it processes health data, privacy laws such as GDPR (or HIPAA in the US) apply, usually triggering a DPIA, and sector-specific requirements such as clinical validation or FDA/EMA guidance come on top. Overlooking these steps could block commercialization entirely, in addition to fines and patient-safety liability.
- **An AI-powered hiring platform:** A resume-screening system falls under the high-risk categories of the EU AI Act and Colorado's AI Act (as a substantial factor in a consequential decision), as well as specific rules such as New York City's Local Law 144, which mandates bias audits for automated hiring tools. The organization would need impact assessments, candidate notifications, bias testing, and human review of

automated rejections. Ignoring these obligations exposes the company to discrimination claims and regulatory penalties.

- **A customer support chatbot:** A generative AI assistant is generally a limited-risk system under the EU AI Act, mainly subject to transparency obligations (users must know they are interacting with a machine), while privacy laws such as GDPR or CCPA/CPRA govern any personal data it collects. In China, additional registration and content moderation duties would apply. The compliance effort is lighter, but a chatbot that mishandles personal data or gives harmful advice can still create legal and reputational damage.

As you can see, the applicable mix depends on what the system does, where it operates, and which sector it touches, and the compliance effort scales with the level of risk. The following comparative table summarizes the major regulations covered in this section, as shown in *Table 9.2*:

| Regulation                                                     | Scope       | Geography      | Primary objective                                                | Key obligations                                                                                                                                                                  |
|----------------------------------------------------------------|-------------|----------------|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EU AI Act                                                      | AI-specific | European Union | Risk-based regulation of AI systems                              | Risk management, data quality, transparency, and human oversight for high-risk systems                                                                                           |
| Colorado AI Act (SB 26-189, repealing and replacing SB 24-205) | AI-specific | Colorado (US)  | Preventing algorithmic discrimination in consequential decisions | Developer documentation, consumer notice and plain-language disclosure, right to correction and human review, AG enforcement with a 60-day cure period, effective 1 January 2027 |

|                                                                         |                 |                 |                                                    |                                                                                            |
|-------------------------------------------------------------------------|-----------------|-----------------|----------------------------------------------------|--------------------------------------------------------------------------------------------|
| China's AI framework                                                    | AI-specific     | China           | Oversight of recommender and generative AI systems | Registration, content moderation, and ethics assessments                                   |
| Canada's AIDA (lapsed as Bill C-27 died in January 2025)                | AI-specific     | Canada          | Regulation of high-impact AI systems               | Risk and impact assessments, bias mitigation, transparency, and record-keeping             |
| GDPR                                                                    | Data protection | European Union  | Protection of personal data and individual rights  | Lawful processing, consent, individual rights, and DPIAs                                   |
| CCPA/CPRA                                                               | Data protection | California (US) | Consumer privacy rights                            | Access, deletion, and opt-out rights, plus rules on sensitive data and automated decisions |
| Sector-specific rules (e.g., HIPAA, FCRA, ISO 26262, NYC Local Law 144) | Vertical        | Varies          | Safety and fairness in critical industries         | Clinical validation, audits, bias testing, and certifications                              |

Table 9.2: A comparative summary of major AI-related regulations

Precisely this last item connects with the next section, in which you will learn more about relevant ISO standards for artificial intelligence and their relation with AI and other regulations.

The rising role of AI standards

Think about this. The ethical principles explain the "what", and the regulations try to connect that "what" with the "why", or the reason behind those regulations and their principles. Now, as AI adopters, you and your organization try to find the "how" to see how to implement this ethical and responsible AI approach. But that's the reality for most of the organizations around the



world, and that is why international standards help define some frameworks and best practices to achieve your goals. Keep in mind that standards are not mandatory, while regulations are.

As with any other technology area, international organizations like the **International Organization for Standardization (ISO)** or the **European Committee for Standardization (CEN)** and the **European Committee for Electrotechnical Standardization (CENELEC)** approach standardization based on a series of technical standards and reports. From an ISO point of view, here is a selection of the main topics you need to keep in mind (extracted directly from the information in *Chapter 1*):

- **ISO 42001:** This standard focuses on AI management systems. It outlines the requirements for establishing, implementing, maintaining, and continually improving an AI management system within an organization.
- **ISO 22989:** This standard provides concepts and terminology for artificial intelligence (AI). It establishes a common understanding of AI terms and definitions to facilitate communication and collaboration across different sectors.
- **ISO 23894:** This standard offers a guide for AI risk management. It provides principles and guidelines for identifying, assessing, and managing risks associated with AI technologies.
- **ISO Guide 73:** This guide defines risk management terminology. It aims to create a consistent understanding of risk management terms and concepts across different industries and sectors.
- **ISO 31000:** This standard provides guidelines for risk management. It offers a framework and process for managing risks that can be applied to any organization, regardless of its size, industry, or sector.
- **ISO 31010:** This standard describes various risk assessment techniques. It provides guidance on selecting and applying techniques to identify and assess risks in different contexts.
- **ISO 38500:** This standard focuses on IT governance. It provides principles, definitions, and a model for effective governance of IT within organizations.
- **ISO 27000:** This standard is part of the ISO/IEC 27000 family of standards, which provides an overview and introduction to information security management systems (ISMS). It includes terms and definitions used throughout the ISMS family of standards.

- **ISO 5259-1:** This standard addresses data quality for analytics and machine learning (ML). It provides guidelines for ensuring the quality of data used in analytics and ML processes.
- **ISO TR 24028:** This technical report offers a framework for AI trustworthiness. It provides guidelines for assessing and enhancing the trustworthiness of AI systems, focusing on aspects such as transparency, accountability, and reliability.

These and other standards collectively cover a wide range of topics related to AI, risk management, IT governance, and information security, providing comprehensive guidelines and frameworks for organizations to follow.

Other relevant international standardization initiatives include the CEN and CENELEC **Joint Technical Committee 21 (JTC 21)**, dedicated to developing European standards for Artificial Intelligence, including harmonization in support of the EU AI Act, or the NIST AI guidelines in the United States of America.

NIST's work on AI, particularly the **NIST AI Risk Management Framework (AI RMF)**, is much more detailed and practice-oriented than many other standards. The AI RMF provides a structured approach to identifying, assessing, and managing risks across the AI lifecycle. It breaks down risk management into clear functions (e.g., govern, map, measure, and manage) and provides concrete recommendations for implementing trustworthy AI, including considerations like explainability, robustness, and bias mitigation. NIST's guidelines are voluntary and allow quick iterations based on new research and feedback.

All of these standards operate under certifiable management-system standards such as ISO/IEC 42001, which operate under a plan-do-check-act cycle backed by internal and external audits, whereas voluntary frameworks such as the NIST AI RMF rely on self-assessment.

## Notion of data and AI governance

Data governance and AI governance are related but very different domains, each focusing on different layers of control and oversight. AI governance relies on data governance as a prerequisite or baseline good practice, but it introduces additional layers of complexity, oversight, and societal accountability. Concretely:

- **Data governance:** It refers to the management of data assets through policies, standards, and procedures that ensure data is accurate, consistent, secure, and used appropriately throughout its lifecycle. It focuses on defining roles and responsibilities for data stewardship, maintaining data quality, ensuring regulatory compliance, and managing data access and ownership. The primary objective is to ensure that data, as a foundational resource, is trustworthy and usable across an organization.

Compared to AI governance, data governance is a way more mature domain of work with existing bodies of knowledge, frameworks, and standards. Some of the most relevant ones are:

- **DAMA-Data Management Body of Knowledge (DMBOK):** A comprehensive framework developed by the **Data Management Association (DAMA)** that defines eleven core data management knowledge areas, including data governance, data quality, metadata, and data architecture, emphasizing roles, policies, and best practices across the data lifecycle.
- **Data Management Capability Assessment Model (DCAM):** Created by the EDM Council, DCAM is a maturity model that provides structured guidance for evaluating and improving enterprise data management capabilities, with a focus on governance, data quality, analytics, and risk management, especially in financial services.
- **Cloud Data Management Capabilities (CDMC):** Also, from the EDM Council, CDMC extends DCAM principles to cloud environments, offering best practices and benchmarks for securely managing sensitive data in hybrid and multi-cloud infrastructures, with a strong focus on security, privacy, and compliance.
- **Control Objectives for Information and Related Technologies (COBIT):** A governance and management framework from ISACA designed for enterprise IT, including data governance. COBIT provides principles, objectives, and enablers to align IT strategies with business goals, emphasizing control, compliance, and performance monitoring.
- **ISO/IEC 38505 for "Governance of Data":** A companion to ISO 38500 (IT Governance), this standard focuses on governance of data as a strategic asset, providing high-level principles for accountability, stewardship, and value delivery in data usage across enterprises.
- **FAIR Data Principles:** Concretely, a set of guidelines to improve data stewardship in research and academia, focused on enhancing data sharing and reuse by both humans and machines, increasingly influential in open science and government initiatives. FAIR stands for Findable, Accessible, Interoperable, and Reusable, as positive attributes for data assets.
- **AI governance:** Compared to data governance, it operates at a higher level of abstraction and complexity. It encompasses the entire lifecycle of artificial intelligence systems and introduces a broader range of oversight mechanisms. While it depends on solid data governance, AI governance extends to how algorithms are designed, trained, validated, deployed, and monitored. It addresses ethical and operational issues such as

model accountability, bias mitigation, fairness, safety, transparency, and explainability. AI governance involves setting standards for responsible innovation, assessing and managing risks, and establishing controls to ensure that AI systems align with organizational values, legal norms, and public expectations.

To see how the two disciplines complement each other in practice, consider an AI-powered loan approval system as a running example. Data governance ensures that the applicant data feeding the model is accurate, complete, up to date, and lawfully collected, that access to sensitive financial information is restricted and logged, and that clear owners are accountable for data quality issues. AI governance then builds on that foundation: it requires the credit-scoring model to be tested for fairness across demographic groups, its decision logic to be explainable to both regulators and rejected applicants, a human reviewer to be able to override automated rejections, and its performance and drift to be monitored continuously after deployment. If the underlying data governance fails, no amount of model-level oversight will save the system; if AI governance is missing, even perfect data can produce unfair or opaque decisions.

AI governance can be approached through a combination of top-down and bottom-up strategies, each offering distinct pathways for implementing oversight, ensuring accountability, and promoting responsible AI development and deployment within the organization. Concretely:

- **Top-down approach:** Top-down AI governance is driven by executive mandates, formal policies, and regulatory compliance frameworks. It is characterized by centralized decision-making, where rules, standards, and oversight mechanisms are defined by leadership or governing bodies and implemented across the organization.

This approach typically involves the establishment of corporate-level AI ethics boards, governance councils, or dedicated Responsible AI offices. These entities are responsible for issuing organization-wide policies, conducting risk assessments, setting approval thresholds for high-risk AI systems, and ensuring compliance with external laws and standards. Audits, formal documentation requirements (like model cards or algorithmic impact assessments), and mandatory checkpoints at different stages of the AI lifecycle are common governance mechanisms under this model.

Top-down governance ensures consistency, enforces accountability at the strategic level, and aligns AI initiatives with corporate values and regulatory obligations. However, it can become rigid or disconnected from the realities of AI development if it lacks regular feedback from operational teams.

- **Bottom-up approach:** Bottom-up AI governance starts from technical and operational teams, such as data scientists, engineers, and product managers. These team members adopt and champion ethical practices and risk mitigation strategies during the actual development and deployment of AI systems. It is driven by practical knowledge, experimentation, and iterative learning rather than centralized policy.

In this model, AI practitioners integrate methods such as fairness toolkits, interpretability libraries, model documentation practices, and human-in-the-loop design patterns into their daily workflows. Localized ethics reviews, peer-model evaluations, and participatory design practices with affected users also represent grassroots methods of governance. Additionally, team-level innovation can identify risks early and escalate issues to higher governance bodies.

Bottom-up governance encourages ownership, agility, and innovation. It ensures that ethical and safety considerations are embedded directly into development practices. However, without alignment to higher-level goals and oversight, it risks fragmentation, inconsistent standards, or blind spots in risk evaluation.

- **Integrated approach:** A mature AI governance strategy combines both approaches through a multi-tiered escalation model. In such a system, lower-level teams have the autonomy to implement governance practices and flag issues, while clear pathways exist to escalate concerns to centralized oversight bodies for resolution. For example, a technical team might raise a concern about fairness during model validation, triggering a review by a Responsible AI Council or Ethics Committee before deployment proceeds. This integrated model balances control and team involvement to encourage individual and group accountability during the AI governance exercise.

There is no good or bad way to implement AI governance and the related escalation processes. Your role as an AI leader in your organization focuses on facilitating the discussions with your executive, operational, and technical teams to guarantee you find a suitable approach that brings actual value to the company. And of course, the scope of the governance process will depend on your own mix of regulations, principles, standards, and any other strategic goals and best practices.

## Role of the RAI Champions

As much as you want to implement a proper AI governance system in your company, you will have to rely on specific individuals who can help implement all processes during the day-to-day activities, and who serve as a bridge between the C-suite and the rest of the team members.

For that purpose, the **Responsible AI (RAI) Champion** plays a pivotal role in making AI governance tangible by translating high-level principles such as fairness, transparency, accountability, or safety into actionable practices that are embedded throughout an organization's AI lifecycle. Unlike purely symbolic roles or abstract committees, the RAI Champion serves as both a strategist and an enabler, ensuring that ethical AI is not just a policy aspiration but a concrete part of everyday development and deployment processes. Keep in mind that being an RAI Champion doesn't mean holding an official role in your organization, but a functional responsibility that can be compatible with the main job as project manager, product owner, or any other.

This means that the RAI Champion acts as a central coordinator who aligns diverse stakeholders around a unified AI governance framework. This involves developing internal policies, interpreting external regulations and standards, and creating tailored governance structures that fit the organization's culture and maturity. By contextualizing requirements within specific use cases, the RAI Champion helps teams operationalize the main concepts.

One of the most tangible contributions of the RAI Champion is the implementation of practical tools and workflows that integrate governance into existing AI development pipelines, but also the facilitation of education activities and cultural change. For example, an RAI Champion can lead training sessions, facilitate workshops, and build communities of practice that empower employees to engage with ethical AI considerations. They can also serve as an escalation point for unresolved dilemmas, based on the escalation approaches you have learned about in previous sections.

But just as no AI governance system is implemented without the employees' participation, you cannot expect tangible results without technical tools that enable and automate certain processes, policies, and rules.

## Illustrative toolkit to implement RAI

A proper RAI toolkit includes technologies, processes, and documentation that collectively embed accountability, safety, fairness, and human oversight into AI systems. This toolkit supports a wide range of activities related to the entire AI lifecycle. The toolkit can be organized into three interdependent categories:

1. **Tools for data governance:** As you already know, data governance tools focus on the quality, integrity, provenance, and compliance of data used in AI systems. These are foundational to any responsible AI program, as poor data governance undermines fairness, reliability, and legal compliance.
  - **Data catalogs and lineage tools:** Help track where data comes from, how it changes over time, and who is responsible for it.

- **Metadata management systems:** Ensure that data attributes, collection context, and usage limitations are documented, enabling informed use of data in modeling.
  - **Privacy engineering tools:** Implement techniques like differential privacy, anonymization, or secure access control for compliant data usage.
  - **Data quality monitoring:** Platforms like Talend or Great Expectations automatically detect anomalies, missing values, or schema violations in datasets used for training AI models.
2. **Tools for AI governance:** These tools orchestrate risk management, accountability, and compliance across the lifecycle of AI systems. They provide central visibility into which AI models exist, what risks they carry, and how they are being managed.
- **AI governance platforms:** Centralize policies, track compliance, score models for risk, and manage audits and approval workflows for AI deployment.
  - **Algorithmic risk classification tools:** Assign risk tiers (e.g., low, medium, high) based on model context, impact, and autonomy, helping prioritize oversight.
  - **Automated documentation systems:** Create audit-ready logs and reports of model performance, drift, bias, and decision logic.
  - **Red teaming and adversarial testing suites:** Stress-test AI systems against edge cases, manipulation, or bias propagation under real-world or adversarial conditions.
3. **Other technical tools:** Besides regular governance tools, the technical toolkits will help you implement specific ethical AI principles, directly connected to concrete use cases and AI implementations:
- **Explainability:** These tools allow stakeholders to understand how and why an AI system makes decisions. These tools support transparency, accountability, and compliance with laws requiring meaningful explanation.
    - **SHapley Additive exPlanations (SHAP):** Provides local and global explanations by assigning importance scores to input features.
    - **Local Interpretable Model-agnostic Explanations (LIME):** Generates approximations of complex models for interpretability at the instance level.
    - **Interpretable models:** Use inherently transparent algorithms (e.g., decision trees, linear models) in regulated or high-risk contexts when feasible.

- **Causal inference tools:** Platforms that go beyond correlations to understand causal relationships, often useful in sensitive applications like healthcare or social services.
- **Explainability dashboards:** Visualize how inputs affect outputs, often combined with fairness assessments to aid governance reviews.
- **Transparency:** Transparency tools enhance visibility into the development, use, and limitations of AI systems. They foster trust, facilitate compliance, and support ethical accountability.
  - **Model cards:** Structured summaries that document a model's purpose, performance, training data, ethical considerations, and known limitations. These are readable by both technical and non-technical stakeholders and are becoming industry standard.
  - **Datasheets for datasets:** Provide detailed metadata and context for datasets, such as how the data was collected, its structure, usage constraints, and potential biases.
  - **AI Software Bill of Materials (AI-SBOM):** Inspired by cybersecurity practice, this tool lists all components, including open-source libraries, datasets, and APIs, used in an AI system. It enables traceability, reproducibility, and vulnerability tracking in the AI supply chain.
  - **Agent cards:** Designed for autonomous or interactive agents (e.g., LLM-based assistants), these would describe agent identity, behavior policies, learning mechanisms, and interaction history. They will be key for tracking accountability in increasingly autonomous systems.

This is a highly evolving space with new functional and technical needs, and an increasing list of providers and tools. Keep track of any new development or tools that may be useful for you and your company.

## Turning insights into action

The reading is done, but the real work starts now. To ensure this playbook doesn't just sit on a shelf, here is your actionable blueprint to translate everything you've learned into momentum.

Each action references the chapter or chapters where the relevant framework or template was introduced, so you can revisit that material for a quick refresher before executing the step, without needing to re-read the chapters in full:



**Note****Your 90-Day AI Leadership Blueprint**

As a top-down AI leader, your next steps must be deliberate. Use this 90-day phased approach to operationalize the frameworks, templates, and strategies covered in this book.

**Days 1–30: Set the Vision, Strategy, and Guardrails**

- ☐ Before dedicating more resources to your AI initiative or signing a vendor contract, you must establish the rules of the game.
- ☐ Publish the AI Policy (*Chapter 1*): Adapt the Company AI Usage Policy template and distribute it. Give your organization immediate clarity on approved tools, prohibited use cases, and data handling requirements.
- ☐ Establish the AI Committee (*Chapter 1 and 9*): Form your cross-functional governance body. Ensure you have the right mix of business, technology, legal, and HR leadership.
- ☐ Define Your AI North Star (*Chapter 1 and 5*): Revisit Porter's Generic Strategies. Are you using AI for cost leadership, differentiation, or focus? Ensure your executive team is aligned on why you are investing in AI.
- ☐ Assess Your AI Maturity (*Chapter 5*): Use the maturity assessment to honestly evaluate whether your organization is Exploring, Planning, Implementing, Scaling, or Realizing. Map the gaps in your business, tech, and cultural pillars.

**Add your notes here:**

---

---

---

---

---

**Days 31–60: Align Use Cases, Talent, and Technology**

- ☐ With the guardrails in place, it is time to evaluate your opportunities and build the engine that will deliver them.
- ☐ Run a Use Case Discovery Workshop (*Chapters 3 and 4*): Bring your business units together. Use the ML Canvas (*Chapter 3*) and the AI capabilities list (Forecasting, Classification, Detection, Completion, Content creation, Automation, Optimization), together with the generative AI and RAG patterns from *Chapter 4*.

☐ **Design Your AI Operating Model (Chapter 7):** Decide how you will structure your teams. Will you use a centralized Center of Excellence, decentralized pods, or a hybrid hub-and-spoke model? Clarify ownership so that product, engineering, and compliance teams don't bottleneck each other.

☐ **Audit Your Talent and Reskilling Needs (Chapter 7):** Evaluate your current workforce. Do you need a Generative AI Architect? Are your Data Scientists adapting to MLOps/LLMOps? Start building a culture of psychological safety where bottom-up innovation, "vibe-coding", and experimentation are rewarded, but controlled.

☐ **Evaluate Your Vendor Ecosystem (Chapter 8):** Use the Vendor Checklist to assess your current and proposed AI tools. Ensure your cloud infrastructure provides the scalability you need while maintaining an architectural layer that prevents hard vendor lock-in.

**Add your notes here:**

---

---

---

---

---

### **Days 61–90: Execute, Measure, and Govern**

☐ **Transition from planning to disciplined execution,** ensuring that your AI investments are sustainable, responsible, and driving real value.

☐ **Implement FinOps for AI (Chapter 6):** AI costs can spiral out of control if left unmonitored. Define your token economics and set up your pricing, chargeback, and ROI measurement models. Move beyond basic "automation" metrics to track actual business transformation.

☐ **Appoint Responsible AI Champions (Chapter 9):** Decentralize your governance by naming RAI champions within your delivery pods. Equip them with the tools to perform fairness audits, data lineage tracking, and bias reduction.

☐ **Launch Your Strategic AI Roadmap (Chapter 5):** Consolidate your prioritized use cases into a 1-year tactical roadmap and a 3-year strategic vision. Schedule the necessary quarterly reviews and communicate this roadmap transparently to investors, board members, and employees.

**Add your notes here:**

---

---

---

---

---

As you navigate the next few months of AI leadership, remember the lessons of past hype cycles. The market will continue to evolve at a crazy pace. New models will be released, new startups will promise the moon, and the regulatory landscape will continue to shift.

To survive and thrive, you must remain the anchor for your organization. Follow these three principles:

1. **Don't fall for AI-washing.** Demand clear use cases, measurable outcomes, and rigorous vendor evaluation.
2. **Protect your people.** The shift to an AI-first organization creates anxiety. Be the leader who communicates clearly about the future of work, upskilling, and how AI will augment rather than simply replace your workforce.
3. **Always remember that AI governance is an enabler, not a blocker.** Treat security, ethics, and compliance as product features that build customer trust, not as afterthoughts.

## Summary

You made it to the end. From the early dreams of the 1950s Dartmouth Conference to the hyper-accelerated era of Large Language Models and Agentic AI, you have traversed the complex landscape of artificial intelligence. If you felt a touch of "impostor syndrome" when you first picked up this playbook, you should now feel equipped, grounded, and ready to lead.

You now know that leading an AI initiative is not a mere delegation exercise. It is not an IT project. It is a fundamental business transformation. You have explored the technical capabilities of both traditional machine learning and modern generative AI, learned how to design AI-first business models, understood the nuances of the AI technology stack, and recognized that without robust governance and the right talent, even the best algorithms will fail.

In this final chapter, you learned how to move from AI ethics to actual accountability, how to quantify Responsible AI initiatives with concrete metrics and tools, how ethics relate to the fast-evolving landscape of AI regulations and standards, and how to implement data and AI governance with the support of RAI Champions and a practical toolkit.

You now have the vocabulary, the frameworks, and the templates to connect board-level strategy with technical execution. Keep this playbook close, revisit the checklists whenever you launch a new initiative, and embrace the continuous learning loop.

And remember, fortunately the era of generative AI is not just about intelligent machines; it is about intelligent leadership and human talent. It's your time to lead. See you in the future!

Adrián González Sánchez,  
*Your friendly neighborhood author*

**P.S. Feel free to contact me if you need any help with your AI initiatives:**

<https://www.linkedin.com/in/adriangs86>

Note

**End of book notes: Add your additional notes and thoughts here:**

---

---

---

---

---

---

---

---

---

---

---

---

---

---

---



# 10

## Unlock Your Exclusive Benefits

Your copy of this book includes the following exclusive benefits:



### DRM-Free PDF Version

Download DRM-free PDF and ePub copies of this book.



### 7-Day Packt Library Access

Get 7-day unlimited access to 8,000+ books and videos. No credit card required.

Available for first-time Packt+ trial users only.



### Next-Gen Reader Access

Read this book on Packt Reader with progress sync, dark mode and note-taking.

Follow the guide below to unlock them. The process takes only a few minutes and needs to be completed once.

## Unlock this Book's Free Benefits in 3 Easy Steps

### Step 1

Keep your purchase invoice ready for *Step 3*. If you have a physical copy, scan it using your phone and save it as a PDF, JPG, or PNG.

For more help on finding your invoice, visit <https://www.packtpub.com/en-us/unlock?step=1>.

#### Note

**Note:** If you bought this book directly from Packt, no invoice is required. After *Step 2*, you can access your exclusive content right away.

### Step 2


Scan the QR code or go to [packtpub.com/unlock](https://packtpub.com/unlock).



On the page that opens (similar to *Figure 10.1* on desktop), search for this book by name and select the correct edition.

### Unlock Your Book's Free Benefits

Bought a Packt book from Amazon or one of our channel partners? Unlock your free benefits in 3 easy steps.



The image shows a desktop view of the Packt unlock landing page. At the top, there's a progress bar with three steps: 'Find Your Book' (active, orange circle), 'Sign Up or Sign In' (inactive, grey circle), and 'Upload Purchase Proof' (inactive, grey circle). Below the progress bar, there's a 'Need Help?' link. The main content area is divided into three sections: 1. Find Your Book (with a search bar and a purple arrow pointing to it with the text 'Search for your book here'), 2. Sign up (Free) or Sign In, and 3. Upload Purchase Proof. Each section has a chevron icon on the right side.

Figure 10.1: Packt unlock landing page on desktop

## Step 3

After selecting your book, sign in to your Packt account or create one for free. Then upload your invoice (PDF, PNG, or JPG, up to 10 MB). Follow the on-screen instructions to finish the process.

## Need Help

If you get stuck and need help, visit <https://www.packtpub.com/unlock-benefits/help> for a detailed FAQ on how to find your invoices and more. This QR code will take you to the help page.



### Note

**Note:** If you are still facing issues, reach out to [customer@packt.com](mailto:customer@packt.com).







packtpub.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.


## Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

At [www.packtpub.com](http://www.packtpub.com), you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

# Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:




## The Generative AI Career Masterplan

Dr. Ali Arsanjani, Sadid Hasan, Maxime Labonne, Andreas Horn, Leonid Kuligin

ISBN: 978-1-80669-145-6

- Understand the global impact of Generative and Agentic AI on jobs and industries
- Master key GenAI Agentic AI concepts for any professional background
- Identify your best-fit AI career path and transferable skills
- Build familiarity with LlamaIndex, Haystack, Hugging Face, and core RAG workflows
- Build an AI-enhanced personal brand using the 5-Layer LinkedIn Strategy and Portfolio Showcase Framework
- Apply ethical and responsible AI principles and governance in real-world practice
- Create a lifelong learning and upskilling roadmap for sustained growth



## The Profitable AI Advantage

Tobias Zwingmann

ISBN: 978-1-83620-589-0

- Identify and leverage AI opportunities in your business
- Create actionable AI roadmaps and implement them incrementally
- Scale AI operations to enhance business functions
- Design and evaluate AI use cases through feasibility assessments
- Recognize AI trends to quickly prototype and iterate AI solutions
- Manage AI projects and teams effectively without deep technical skills
- Understand fundamental concepts of modern AI, including Generative AI, LLMs, and predictive models

## Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit [authors.packt.com](https://authors.packt.com) and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

## Share your thoughts

Once you've read *Generative AI for Leaders Playbook*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback.



<https://packt.link/r/1806388499>

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

# Index

## Symbols

**1-bit quantization** 144

## A

**AI Committees** 209

**AI Management System (AIMS)** 33

**AI Platform Team** 216

**AI Product Pods** 215

**AI Red Teaming** 141

**AI Risk Management Framework (AI RMF)** 263

**AI Software Bill of Materials (AI-SBOM)** 269

**AI adoption perspective**  
key milestones 60

**AI agents** 121

**AI committee checklist** 17

**AI cost structure**  
data costs 190  
human costs 189  
platform costs 189  
risk and compliance costs 191

**AI cost units**  
managing 188

**AI development cycle**  
evolution of roles 202  
new roles 206  
traditional roles 203

**AI ethics**  
migrating, to accountability 244, 245

**AI governance** 264

**AI leader**  
best practices 198  
feedback loops, enabling 200  
for AI initiative  
pragmatic interview 199  
processes, designing

qualities 199  
role 198

**AI learning** 87  
data splitting 89, 90  
pattern types 88

**AI learning pattern types**  
reinforcement learning 88  
self-supervised learning 88  
supervised learning 88  
unsupervised learning 88

**AI levels** 78  
Artificial General Intelligence (AGI) 79  
Artificial Superintelligence (ASI) 80  
current status 80, 81  
narrow AI (NAI) 78, 79

**AI lifecycle** 90 – 92

**AI management levels** 30  
evaluating 35, 41  
improving, via 32 – 34  
management systems  
Project Management Office (PMO) 31  
strategic AI management 31  
tactical AI management 31  
technical management 31

**AI maturity, organization**  
assessing 158, 159, 162

**AI model training phases** 134, 135  
customization 142  
inference 142  
optimization 143  
post-training 139 – 141  
pre-training 135 – 138

**AI models** 125

**AI platform** 125

**AI regulation** 208

**AI standards**  
rising role 261 – 263

|                                         |                  |                                     |                |
|-----------------------------------------|------------------|-------------------------------------|----------------|
| <b>AI strategy</b>                      |                  | becoming                            | 179 – 182      |
| business strategy                       | 154, 155         | <b>AI-specific regulations</b>      | <b>253</b>     |
| cloud technologies,                     | 234 – 238        | <b>API management</b>               | <b>125</b>     |
| leveraging                              |                  | <b>Accountable AI</b>               | <b>245</b>     |
| internal governance                     | 156              | <b>Agentic AI</b>                   | <b>121</b>     |
| organization and culture                | 155              | <b>AlphaGo</b>                      |                |
| overall maturity                        | 157              | reference link                      | 105            |
| pillars                                 | 154              | <b>accountability</b>               | <b>247</b>     |
| technology strategy                     | 156              | <b>app modernization</b>            | <b>120</b>     |
| turning, into operational               | 165, 166         | <b>artificial intelligence (AI)</b> |                |
| actions                                 |                  | automation                          | 84             |
| <b>AI talent strategy</b>               |                  | classification                      | 82             |
| considerations                          | 209              | completion                          | 83             |
| fundamental, versus                     | 210              | content creation                    | 83             |
| applied research                        |                  | defining                            | 4              |
| skills and responsibilities,            | 211              | detection                           | 82             |
| focusing on                             |                  | forecasting                         | 82             |
| trends, avoiding                        | 210              | historical perspective              | 44 – 57        |
| <b>AI teams</b>                         |                  | lessons, for business               | 59, 60         |
| empowering                              | 202              | main capabilities                   | 81             |
| scaling, key considerations             | 212              | optimization                        | 84, 87         |
| <b>AI teams scaling, considerations</b> |                  | <b>attention mechanism</b>          | <b>132</b>     |
| AI layoff trap, avoiding                | 214              | <b>augmentation</b>                 | <b>213</b>     |
| current workforce,                      | 213              | <b>automation</b>                   | <b>84, 213</b> |
| reskilling                              |                  | <b>B</b>                            |                |
| noise, ignoring                         | 214              | <b>Bag-of-Words (BoW)</b>           | <b>133</b>     |
| talent gaps and planning                | 212              | <b>Bayesian Linear</b>              | <b>97</b>      |
| analysis                                |                  | <b>regression</b>                   |                |
| work issues future,                     | 213              | <b>Bidirectional Encoder</b>        | <b>134</b>     |
| addressing                              |                  | <b>Representations from</b>         |                |
| <b>AI technology stack</b>              | <b>220, 221</b>  | <b>Transformers (BERT)</b>          |                |
| AI infrastructure                       | 222              | <b>Boosted Decision Tree</b>        | <b>97</b>      |
| development studios                     | 223              | <b>regression</b>                   |                |
| governance, compliance                  | 223              | <b>Business Model Canvas</b>        | <b>111</b>     |
| and safety                              |                  | <b>(BMC)</b>                        |                |
| libraries and frameworks                | 223              | <b>backend software</b>             | <b>204</b>     |
| platforms                               | 222              | <b>developer</b>                    |                |
| <b>AI tooling selection criteria</b>    |                  | <b>backpropagation</b>              | <b>102</b>     |
| defining                                | 224 – 234        | <b>batch parallelism</b>            | <b>143</b>     |
| <b>AI transformation levels</b>         | <b>174</b>       | <b>beam search</b>                  | <b>144</b>     |
| business, transforming                  | 176              | <b>optimization</b>                 |                |
| capabilities, augmenting                | 175              | <b>bias reduction</b>               | <b>247</b>     |
| efficiencies, finding                   | 175              | <b>binary classification</b>        | <b>82</b>      |
| ventures, creating                      | 176              | Naïve Bayes                         | 98             |
| <b>AI washing trap</b>                  | <b>186 – 188</b> |                                     |                |
| <b>AI-enabled services</b>              |                  |                                     |                |
| implementing                            | 178, 179         |                                     |                |
| <b>AI-first organization</b>            |                  |                                     |                |

|                                                                                       |                     |                                                           |               |
|---------------------------------------------------------------------------------------|---------------------|-----------------------------------------------------------|---------------|
| two-class Bayes Point Machine                                                         | 98                  | chunking                                                  | 142           |
| two-class Boosted Decision Tree                                                       | 99                  | hybrid search                                             | 142           |
| two-class Decision Forest                                                             | 98                  | reranking                                                 | 142           |
| two-class Decision Jungle                                                             | 99                  |                                                           |               |
| two-class Logistic Regression                                                         | 98                  |                                                           |               |
| two-class SVM                                                                         | 98                  |                                                           |               |
| <b>bottom-up approach</b>                                                             | <b>266</b>          |                                                           |               |
| <b>C</b>                                                                              |                     | <b>D</b>                                                  |               |
| <b>California's Consumer Privacy Act/Privacy Rights Act (CCPA/CPRA)</b>               | <b>257</b>          | <b>DAMA-Data Management Body of Knowledge (DMBOK)</b>     | <b>264</b>    |
| <b>Canada's Personal Information Protection and Electronic Documents Act (PIPEDA)</b> | <b>257</b>          | <b>DL models</b>                                          | <b>102</b>    |
| <b>Center of Excellence (CoE)</b>                                                     | <b>31, 209</b>      | key considerations                                        | 103           |
| <b>China's AI Regulation Framework</b>                                                | <b>255</b>          | types                                                     | 104 – 106     |
| <b>Cloud Data Management Capabilities (CDMC)</b>                                      | <b>264</b>          | <b>Data Management Association (DAMA)</b>                 | <b>264</b>    |
| <b>Colorado AI Act</b>                                                                | <b>254</b>          | <b>Data Management Capability Assessment Model (DCAM)</b> | <b>264</b>    |
| <b>Control Objectives for Information and Related Technologies (COBIT)</b>            | <b>264</b>          | <b>Data Protection Impact Assessments (DPIAs)</b>         | <b>258</b>    |
| <b>capital expenditure (CAPEX)</b>                                                    | <b>63, 188, 236</b> | <b>Development Operations (DevOps)</b>                    | <b>92</b>     |
| <b>chunking</b>                                                                       | <b>142</b>          | <b>data augmentation</b>                                  | <b>137</b>    |
| <b>classification</b>                                                                 | <b>82</b>           | back-translation                                          | 137           |
| <b>cloud administrator and engineer</b>                                               | <b>204</b>          | noise injection                                           | 137           |
| <b>cloud cost literacy</b>                                                            | <b>192</b>          | synonym replacement                                       | 137           |
| <b>code generation</b>                                                                | <b>120</b>          | synthetic data                                            | 137           |
| <b>completion</b>                                                                     | <b>83</b>           | <b>data engineer</b>                                      | <b>204</b>    |
| <b>container-level optimization</b>                                                   | <b>144</b>          | <b>data governance</b>                                    | <b>263</b>    |
| <b>content adaptation</b>                                                             | <b>120</b>          | <b>data parallelism</b>                                   | <b>138</b>    |
| <b>content creation</b>                                                               | <b>83</b>           | <b>data protection and privacy regulations</b>            | <b>257</b>    |
| <b>contextual embeddings</b>                                                          | <b>136</b>          | <b>data scientists</b>                                    | <b>205</b>    |
| <b>conversational interaction</b>                                                     | <b>120</b>          | <b>data split</b>                                         | <b>89</b>     |
| <b>convolutional neural networks (CNNs)</b>                                           | <b>104</b>          | <b>data split, for AI</b>                                 |               |
| <b>customized inference techniques</b>                                                |                     | test                                                      | 90            |
|                                                                                       |                     | training                                                  | 89            |
|                                                                                       |                     | validation                                                | 90            |
|                                                                                       |                     | <b>deep learning (DL)</b>                                 | <b>55, 94</b> |
|                                                                                       |                     | <b>deep reinforcement learning</b>                        | <b>105</b>    |
|                                                                                       |                     | <b>detection</b>                                          | <b>82</b>     |
|                                                                                       |                     | <b>digital intelligence</b>                               | <b>5</b>      |
|                                                                                       |                     | <b>distillation</b>                                       | <b>141</b>    |
|                                                                                       |                     | <b>distributed training</b>                               | <b>137</b>    |



**E**

|                                                                          |            |
|--------------------------------------------------------------------------|------------|
| <b>European Committee for Electrotechnical Standardization (CENELEC)</b> | <b>262</b> |
| <b>European Committee for Standardization (CEN)</b>                      | <b>262</b> |
| <b>European Union's AI Act</b>                                           | <b>253</b> |
| <b>European Union's General Data Protection Regulation (GDPR)</b>        | <b>257</b> |
| <b>electronic health records (EHRs)</b>                                  | <b>258</b> |
| <b>embedding generation</b>                                              | <b>134</b> |
| <b>embeddings</b>                                                        | <b>132</b> |
| contextual embeddings                                                    | 136        |
| sentence embeddings                                                      | 136        |
| word embeddings                                                          | 136        |
| <b>end-to-end governance</b>                                             | <b>125</b> |
| <b>ethical AI</b>                                                        | <b>244</b> |
| <b>ethical principles</b>                                                |            |
| democratic participation                                                 | 246        |
| equity                                                                   | 246        |
| inclusion and diversity                                                  | 246        |
| protection of privacy and intimacy                                       | 246        |
| prudence                                                                 | 246        |
| respect for autonomy                                                     | 246        |
| responsibility                                                           | 246        |
| solidarity                                                               | 246        |
| sustainable development                                                  | 246        |
| well-being                                                               | 246        |
| <b>ethics</b>                                                            |            |
| and regulations, relation between                                        | 252        |
| <b>explainability</b>                                                    | <b>247</b> |
| <b>external signals and pressures</b>                                    |            |
| handling                                                                 | 67         |

**F**

|                                        |            |
|----------------------------------------|------------|
| <b>FAIR Data Principles</b>            | <b>264</b> |
| <b>Fast Forest Quantile regression</b> | <b>97</b>  |
| <b>Fast Forest regression</b>          | <b>97</b>  |
| <b>FinOps</b>                          | <b>236</b> |

|                                                    |                 |
|----------------------------------------------------|-----------------|
| <b>Forward Deployed Engineer (FDE)</b>             | <b>208</b>      |
| <b>feature engineering</b>                         | <b>90</b>       |
| <b>feedforward neural network (FNN)</b>            | <b>104</b>      |
| <b>financial operations (FinOps)</b>               | <b>192</b>      |
| <b>fine-tuning</b>                                 | <b>128, 139</b> |
| PEFT fine-tuning                                   | 139             |
| regular fine-tuning                                | 139             |
| <b>forecasting</b>                                 | <b>82</b>       |
| <b>frontend software developer</b>                 | <b>203</b>      |
| <b>G</b>                                           |                 |
| <b>Generative Adversarial Networks (GANs)</b>      | <b>134</b>      |
| <b>Generative Pre-trained Transformers (GPT)</b>   | <b>134</b>      |
| <b>Graphics Processing Unit (GPU)</b>              | <b>103, 136</b> |
| <b>generative AI</b>                               | <b>60 – 64</b>  |
| adoption                                           | 126 – 129       |
| architecture                                       | 207             |
| assessing, for innovation and executive engagement | 6, 7            |
| committee, defining                                | 17, 25          |
| customization patterns                             | 126 – 129       |
| enhancements                                       | 118             |
| North Star, confirming for AI use cases            | 25              |
| policy, creating                                   | 7, 8            |
| Porter's Generic Strategies Framework, using       | 25, 26, 29      |
| token economy                                      | 193             |
| <b>generative AI Cloud Expert</b>                  | <b>206</b>      |
| <b>generative AI PM</b>                            | <b>207</b>      |
| <b>generative AI boom</b>                          | <b>5</b>        |
| <b>generative AI developer</b>                     | <b>206</b>      |
| <b>generative AI, building blocks</b>              |                 |
| AI models                                          | 125             |
| AI platform                                        | 125             |
| API management                                     | 125             |
| end-to-end governance                              | 125             |
| guardrails                                         | 125             |
| interfaces                                         | 125             |
| Lakehouse                                          | 124             |

|                                                             |                |                                                               |                       |
|-------------------------------------------------------------|----------------|---------------------------------------------------------------|-----------------------|
| memory                                                      | 124            | <b>interfaces</b>                                             | <b>125</b>            |
| orchestration engine                                        | 125            | <b>interpretability</b>                                       | <b>247</b>            |
| physical infrastructure                                     | 125            |                                                               |                       |
| <b>generative AI, capabilities</b>                          |                | <b>J</b>                                                      |                       |
| access to tools                                             | 121            | <b>Joint Technical Committee 21 (JTC 21)</b>                  | <b>263</b>            |
| app modernization                                           | 120            |                                                               |                       |
| code generation                                             | 120            | <b>K</b>                                                      |                       |
| content adaptation                                          | 120            | <b>K-means</b>                                                | <b>96</b>             |
| conversational interaction                                  | 120            | <b>knowledge management</b>                                   | <b>121</b>            |
| intent detection                                            | 120            |                                                               |                       |
| knowledge management                                        | 121            | <b>L</b>                                                      |                       |
| language understanding                                      | 120            | <b>Lakehouse</b>                                              | <b>124</b>            |
| multimodal generation                                       | 121            | <b>LangSmith</b>                                              | <b>185</b>            |
| multimodal understanding                                    | 121            | <b>Large Language Model Operations (LLMOps)</b>               | <b>93, 237</b>        |
| personalization                                             | 121            | <b>Latent Dirichlet Allocation (LDA)</b>                      | <b>134</b>            |
| reasoning                                                   | 120            | <b>Local Interpretable Model-agnostic Explanations (LIME)</b> | <b>268</b>            |
| semantic search                                             | 120            | <b>Low-Rank Adaptation (LoRA)</b>                             | <b>139</b>            |
| summarization                                               | 120            | <b>language understanding</b>                                 | <b>120</b>            |
| text generation                                             | 120            | <b>large language models (LLMs)</b>                           | <b>5, 56, 93, 118</b> |
| text to query                                               | 121            | <b>linear regression</b>                                      | <b>97</b>             |
| <b>geopolitics</b>                                          | <b>68 – 72</b> | <b>M</b>                                                      |                       |
| <b>guardrails</b>                                           | <b>125</b>     | <b>ML Engineer</b>                                            | <b>205</b>            |
|                                                             |                | <b>ML and DL cases</b>                                        | <b>106, 109 – 111</b> |
| <b>H</b>                                                    |                | <b>ML canvas</b>                                              | <b>111, 112</b>       |
| <b>hybrid search</b>                                        | <b>142</b>     | <b>ML models</b>                                              | <b>95</b>             |
| <b>hype and market signals</b>                              |                | supervised ML models                                          | 96                    |
| turning, into practical advantage                           | 68             | unsupervised ML models                                        | 95                    |
| <b>hype cycles</b>                                          |                | <b>MLflow</b>                                                 | <b>185</b>            |
| lessons learned                                             | 66, 67         | <b>Machine Learning Operations (MLOps)</b>                    | <b>93, 237</b>        |
|                                                             |                | <b>Mixture of Experts (MoE)</b>                               | <b>138</b>            |
| <b>I</b>                                                    |                | <b>Model Context Protocol (MCP)</b>                           | <b>121, 122</b>       |
| <b>Infrastructure as Code (IaC)</b>                         | <b>92</b>      | <b>Model-as-a-Service (MaaS)</b>                              | <b>236</b>            |
| <b>International Organization for Standardization (ISO)</b> | <b>32, 262</b> | <b>machine learning (ML)</b>                                  | <b>94</b>             |
| <b>inclusiveness and fairness</b>                           | <b>247</b>     | <b>memory</b>                                                 | <b>124</b>            |
| <b>inevitable impostor syndrome</b>                         | <b>6</b>       |                                                               |                       |
| <b>inference</b>                                            | <b>142</b>     |                                                               |                       |
| <b>innovation</b>                                           | <b>213</b>     |                                                               |                       |
| <b>insights</b>                                             |                |                                                               |                       |
| turning, into actions                                       | 269, 272       |                                                               |                       |
| <b>integrated approach</b>                                  | <b>266</b>     |                                                               |                       |
| <b>intent detection</b>                                     | <b>120</b>     |                                                               |                       |

|                                                |                     |
|------------------------------------------------|---------------------|
| <b>memory handling</b>                         | <b>143</b>          |
| <b>model parallelism</b>                       | <b>138</b>          |
| <b>multiclass classification</b>               |                     |
| multiclass Decision Forest                     | 99                  |
| multiclass Decision Jungle                     | 99                  |
| multiclass Logistic Regression                 | 99                  |
| One-vs-All Multiclass                          | 99                  |
| <b>multiclass neural network</b>               | <b>104</b>          |
| <b>multimodal generation</b>                   | <b>121</b>          |
| <b>multimodal understanding</b>                | <b>121</b>          |
| <br><b>N</b>                                   |                     |
| <b>Non-negative Matrix Factorization (NMF)</b> | <b>134</b>          |
| <b>natural language processing (NLP)</b>       | <b>120, 132</b>     |
| <b>neural network regression</b>               | <b>104</b>          |
| <b>neural networks</b>                         | <b>102</b>          |
| <b>neurons</b>                                 | <b>102</b>          |
| <br><b>O</b>                                   |                     |
| <b>one-class support vector machine (SVM)</b>  | <b>95</b>           |
| <b>open-source alternatives</b>                | <b>241</b>          |
| <b>operating model</b>                         |                     |
| defining                                       | 214 – 216           |
| <b>operational expenditure (OPEX)</b>          | <b>63, 188, 236</b> |
| <b>optimizations</b>                           | <b>84, 87</b>       |
| 1-bit quantization                             | 144                 |
| batch parallelism                              | 143                 |
| beam search optimization                       | 144                 |
| container-level optimization                   | 144                 |
| memory handling                                | 143                 |
| prompt optimization                            | 144                 |
| semantic caching                               | 143                 |
| top-K sampling                                 | 144                 |
| <b>orchestration engine</b>                    | <b>125</b>          |
| <b>ordinal regression</b>                      | <b>97</b>           |
| <b>organizational scalability</b>              | <b>182 – 184</b>    |

## P

|                                                   |                 |
|---------------------------------------------------|-----------------|
| <b>PEFT fine-tuning</b>                           | <b>139</b>      |
| <b>Parameter-Efficient Fine-Tuning (PEFT)</b>     | <b>139</b>      |
| <b>Poisson regression</b>                         | <b>97</b>       |
| <b>Project Management Office (PMO)</b>            | <b>31</b>       |
| <b>personalization</b>                            | <b>121</b>      |
| <b>physical infrastructure</b>                    | <b>125</b>      |
| <b>pipeline parallelism</b>                       | <b>138</b>      |
| <b>post-training phase techniques</b>             |                 |
| AI Red Teaming                                    | 141             |
| distillation                                      | 141             |
| fine-tuning                                       | 139             |
| pruning                                           | 140             |
| quantization-aware training (QAT)                 | 141             |
| Reinforcement Learning from Human Feedback (RLHF) | 139             |
| <b>pre-training phase techniques</b>              |                 |
| data augmentation                                 | 137             |
| distributed training                              | 137             |
| Mixture of Experts (MoE)                          | 138             |
| multi-modal expansion                             | 137             |
| Self-Supervised Learning (SSL)                    | 135, 136        |
| vectorization                                     | 136             |
| <b>principal component analysis (PCA)</b>         | <b>96</b>       |
| <b>printing press</b>                             | <b>65</b>       |
| <b>prioritization models</b>                      | <b>149, 150</b> |
| <b>product and data analysts</b>                  | <b>205</b>      |
| <b>product development</b>                        |                 |
| impact of AI                                      | 177             |
| <b>prompt engineering</b>                         | <b>127</b>      |
| <b>prompt engineers</b>                           | <b>207</b>      |
| <b>prompt optimization</b>                        | <b>144</b>      |
| <b>pruning</b>                                    | <b>140</b>      |

## Q

|                                          |            |
|------------------------------------------|------------|
| <b>Quality Assurance (QA)</b>            | <b>205</b> |
| <b>quantization-aware training (QAT)</b> | <b>141</b> |

**R**

|                                                          |                     |
|----------------------------------------------------------|---------------------|
| <b>RAI</b>                                               |                     |
| implementing, toolkit                                    | 267 – 269           |
| <b>RAI Champions</b>                                     |                     |
| role                                                     | 266                 |
| <b>RAI initiatives</b>                                   |                     |
| quantifying                                              | 248, 252            |
| <b>Receiver Operating Characteristic (ROC)</b>           | 237                 |
| <b>Responsible AI (RAI)</b>                              | 245                 |
| <b>Responsible AI (RAI) Champion</b>                     | 267                 |
| <b>Responsible AI (RAI) experts</b>                      | 208                 |
| <b>Retrieval Augmented Generation (RAG)</b>              | 127, 136            |
| <b>reasoning</b>                                         | 120                 |
| <b>recurrent neural networks (RNNs)</b>                  | 104                 |
| <b>regression</b>                                        |                     |
| Bayesian Linear regression                               | 97                  |
| Boosted Decision Tree regression                         | 97                  |
| Fast Forest Quantile regression                          | 97                  |
| Fast Forest regression                                   | 97                  |
| linear regression                                        | 97                  |
| ordinal regression                                       | 97                  |
| Poisson regression                                       | 97                  |
| <b>regular fine-tuning</b>                               | 139                 |
| <b>regulations</b>                                       |                     |
| types                                                    | 253, 254, 257 – 260 |
| <b>reinforcement learning</b>                            | 88                  |
| <b>reinforcement learning from human feedback (RLHF)</b> | 88, 139             |
| <b>reranking</b>                                         | 142                 |
| <b>resilience</b>                                        |                     |
| building, while avoiding vendor lock-in                  | 238                 |
| technical flexibility and governance                     | 239                 |
| <b>risk and compliance costs</b>                         | 191                 |

**S**

|                                                           |           |
|-----------------------------------------------------------|-----------|
| <b>SHapley Additive exPlanations (SHAP)</b>               | 268       |
| <b>Self-Supervised Learning (SSL)</b>                     | 135, 136  |
| <b>Subject Matter Experts (SME)</b>                       | 209       |
| <b>scalability</b>                                        |           |
| managing                                                  | 182       |
| organizational scalability                                | 182 – 184 |
| technical scalability                                     | 184, 185  |
| <b>security engineer</b>                                  | 205       |
| <b>self-supervised learning</b>                           | 88        |
| <b>semantic caching</b>                                   | 143       |
| <b>semantic search</b>                                    | 120       |
| <b>sentence embeddings</b>                                | 136       |
| <b>small language models (SLMs)</b>                       | 118       |
| <b>software development lifecycle (SDLC)</b>              | 205       |
| <b>state-of-the-art AI (SOTA)</b>                         | 236       |
| <b>strategic AI roadmap</b>                               |           |
| sharing                                                   | 162, 163  |
| <b>subject matter expert (SME)</b>                        | 95        |
| <b>summarization</b>                                      | 120       |
| <b>supervised ML models</b>                               | 96        |
| classification (binary)                                   | 98, 99    |
| classification (multiclass)                               | 99 – 101  |
| regression                                                | 96 – 98   |
| <b>supervised learning</b>                                | 88        |
| <b>T</b>                                                  |           |
| <b>Taiwan Semiconductor Manufacturing Company (TSMC)</b>  | 71        |
| <b>Tensor Processing Unit (TPU)</b>                       | 103       |
| <b>Term Frequency-Inverse Document Frequency (TF-IDF)</b> | 133       |
| <b>Trustworthy AI</b>                                     | 245       |
| <b>technical operations</b>                               | 92 – 94   |
| <b>technical scalability</b>                              | 184, 185  |

|                                  |            |                                   |                 |
|----------------------------------|------------|-----------------------------------|-----------------|
| <b>text generation</b>           | <b>120</b> | <b>transformer models</b>         | <b>134</b>      |
| <b>text representation</b>       | <b>133</b> | <b>transparency</b>               | <b>246</b>      |
| <b>token literacy</b>            | <b>192</b> | <b>two-class locally deep SVM</b> | <b>104</b>      |
| <b>tokenization</b>              | <b>132</b> | <b>two-class neural network</b>   | <b>104</b>      |
| <b>top-K sampling</b>            | <b>144</b> |                                   |                 |
| <b>top-down approach</b>         | <b>265</b> | <b>U</b>                          |                 |
| <b>topic modeling</b>            | <b>134</b> | <b>UX designers</b>               | <b>205</b>      |
| <b>traditional roles</b>         |            | <b>unsupervised ML models</b>     | <b>95</b>       |
| backend software developer       | 204        | anomaly detection                 | 95, 96          |
| cloud administrator and engineer | 204        | clustering                        | 96              |
| data engineer                    | 204        | <b>unsupervised learning</b>      | <b>88</b>       |
| data scientists and ML engineers | 205        | <b>use case discovery</b>         | <b>146, 147</b> |
| frontend software developer      | 203        | <b>use case evaluation</b>        | <b>148, 149</b> |
| product and data analysts        | 205        |                                   |                 |
| QA and security engineers        | 205        | <b>V</b>                          |                 |
| tactical roles                   | 206        | <b>vectorization</b>              | <b>136</b>      |
| UX designers                     | 205        | <b>vectors</b>                    | <b>132</b>      |
| <b>trainers</b>                  | <b>209</b> | <b>vendor negotiation</b>         | <b>240</b>      |
| <b>training</b>                  | <b>133</b> | <b>W</b>                          |                 |
|                                  |            | <b>word embeddings</b>            | <b>136</b>      |

