ELEMENT AI

FOR MORE INFORMATION
sales@elementai.com

The AI Maturity Framework

A strategic guide to operationalize and scale enterprise AI solutions

WHITEPAPER

TABLE OF CONTENTS

About Element AI 3

Executive Summary 4

Introduction 5

A Framework to Evaluate AI Maturity 7

Introduces the AI Maturity Framework and provides context regarding the state of organizational maturity for AI in industry.
Use this section to orient your thinking about AI maturity overall and decide what to read next.

- The AI Maturity Framework 7
- The State of Organizational AI Maturity 8

The Five Stages of AI Maturity 10

Documents the five stages of AI maturity and how each one unfolds over time, from getting started, to moving forward, to leveling up.
Use this section to gain a deeper understanding of the key challenges and opportunities in your current stage of AI maturity.

- Stage 1: Exploring 12
- Stage 2: Experimenting 13
- Stage 3: Formalizing 14
- Stage 4: Optimizing 15
- Stage 5: Transforming 16

The Five Dimensions of Enterprise AI 17

Details the five dimensions that enable enterprise AI and how each one contributes to advancing AI maturity over time.
Use this section to fine-tune your understanding of each organizational enabler and decide how to prioritize AI efforts to level up.

- Dimension 1: Strategy 19
- Dimension 2: Data 22
- Dimension 3: Technology 26
- Dimension 4: People 29
- Dimension 5: Governance 33

Conclusion: Putting it All Together 36

Summarizes and provides guidance on how to use the framework to advance AI maturity.
Use this section as a quick reference for you and your team to align on how to frame, discover, define, and prioritize next best actions for AI.

The AI Maturity Framework: Executive Blueprint 37

Glossary 38

About
Element AI

Element AI develops AI-powered solutions and services that help people and machines work smarter, together. Founded in 2016 by serial entrepreneurs including JF Gagné and A.M.Turing Award recipient, Yoshua Bengio, PhD, Element AI turns cutting-edge fundamental research into software solutions that exponentially learn and improve. Its end-to-end offering includes advisory services, AI enablement tools and products, aimed at helping large organizations operationalize AI for real business impact. Element AI maintains a strong connection to academia through research collaborations and takes a leadership position in policymaking around the impact of AI technology on society.
www.elementai.com

Executive Summary

Recent progress in artificial intelligence may represent the most significant technological advancement in a generation, but progress is uneven. Our recent industry survey confirms that most enterprise organizations still have not graduated beyond their first AI experiments and pilot projects.

Progress is slow at most enterprises because implementing AI depends on technical as well as organizational factors—and few resources exist to help leaders plan and strengthen their organizational foundations for AI.

In this document, we present a comprehensive AI Maturity Framework to close that gap. The AI Maturity Framework is designed to help leaders understand and prioritize the actions that will have the greatest impact on AI in their unique context. It catalogs five key dimensions that must be aligned to create and scale business impact with AI: Strategy, Data, Technology, People and Governance. It also explains how these dimensions define an organization's maturity across five stages: Exploring, Experimenting, Formalizing, Optimizing and Transforming.

We also address how the AI maturity journey is unfolding across industries today. Throughout the document, we share the firsthand experience of our AI Advisory and Enablement practice as well as provide insights from an industry survey conducted with senior decision-makers between October 2019 and January 2020.

At a macro level, our survey confirms that fewer than one in ten organizations (7%) are mature enough to operationalize and scale AI. About twice as many (14%) are aligning Strategy, Data, Technology, People and Governance to join this vanguard. Another 52% are working through experiments to validate specific business cases for AI.

Our framework, cases and survey data help explain these statistics. We show how mature organizations tend to emphasize Strategy for AI, securing executive sponsorship and clarifying organizational roadmaps early. Many organizations are behind on Governance for AI and still need to set policies and practices for managing new risks. In early stages of maturity, organizations tend to invest in Data for AI before defining data requirements with AI use cases.

Using the framework, and guided by insights from our cases and survey, business leaders can learn how the five organizational dimensions need to evolve in the age of AI, and quickly assess their own progress in each dimension. Then, they can target the best next steps for impact.

Introduction

If you are evaluating, designing, or championing your organization's strategy for using artificial intelligence, this document is for you. It is designed to help senior decision-makers as well as implementation teams, whether you plan to purchase an off-the-shelf AI solution, build one yourself or take a hybrid approach.

We wrote this document because artificial intelligence is animating the world that electricity illuminated and that the Internet connected. But, not unlike electricity in 1910 or the Internet in 1990, AI in 2020 still hasn't made a real impact yet for most businesses. As with any new revolutionary technology, it is taking time for industry leaders to figure out how to leverage it in a tangible and embedded manner.

From our vantage point, we see that while many challenges remain, the tipping point is not far away—and it is closer in some industries than in others. Organizations are discovering that AI is difficult for reasons that go beyond the scientific and technical.

Fundamentally, organizations need to become digital at their core. This is what unlocks the organization's potential to operate without the constraints of traditional enterprises, to compete in new ways, capture unprecedented value and alter the very industries in which it operates. What we are really seeing with AI is a redefinition of what an organization can be—how it operates, strategizes and competes.

What is AI? The goal of the field hasn't fundamentally changed since its inception in the 1950s: to create machines that exhibit human-like intelligence. In seventy-odd years, methods for achieving this goal have proliferated. The field is now a dynamic hybrid of hard science and practical engineering, with dedicated research programs for applications such as machine vision and natural language processing; techniques such as neural networks and reinforcement learning; and social implications such as Fairness, Accountability, and Transparency (FAccT).

Now, AI systems perform at or above human-level for many specialized tasks. This includes tasks that were never before possible or practical to address with written rules or traditional software, such as intelligently recognizing and categorizing millions of images. There are even more creative applications of AI, such as generating new images, text and other data. And fundamental AI research activity is still on the rise.

Yet AI has been difficult for organizations to adopt because organizations have to change how they think, act and learn in order to take advantage of what it offers. And it takes time for organizations to mature their AI capabilities and
the aspects that support AI.

What is AI maturity? It's a measure of an organization's ability to achieve and scale impact from AI systems. Our recent industry survey confirms that in January 2020, fewer than 1 in 10 organizations are mature enough to put AI into production. But about 1 in 7 are actively clarifying their strategy for AI, developing their data and technology infrastructure, aligning their teams, and setting governance practices to scale responsibly.

In early stages, AI maturity typically focuses on improving operations so organizations can achieve their existing strategic goals. For example, optical character recognition (OCR) and natural language processing can streamline document processing so a business can expand its market reach.

In later stages, AI becomes more central to the strategy of the organization itself. Think of the operating models of the FAANG companies or upstart firms like Uber and Grab. AI has broken down silos in these organizations (or silos didn't exist from the start) so human-machine collaboration is free to drive the entire business. At the highest stages of maturity, AI is central to how organizations deliver as well as conceive of new business models, products and services. Cue the emergence of a different kind of a firm with AI as its operating system.

The key to AI maturity, from exploring AI to transforming with it, is envisioning what that end-state could look like for you and envisioning a clear path to that vision from your current state. Most business leaders are behind in being able to grasp either the current or future states clearly. This is the primary driver for why we wrote this document.

This document shares what we've learned from our research and experience as AI practitioners to help you join the vanguard of organizations now using AI—or to get ahead of the pack. The central topic is our detailed framework for assessing AI maturity and focusing on the right actions to level-up. We also include results from our recent survey of senior decision-makers in multiple industries and cases from our advisory practice.

At Element AI, we are inspired by the promise of artificial intelligence. We're also privileged to go on this transformational journey with our clients, to help them realize the promise of AI to create the future of financial services, supply chains, customer experiences, our cities, and our environment.

When we lead our organizations to work smarter with AI, we move the world forward. From illuminated, to connected, to animated—to all that comes next.

Karthik Ramakrishnan
Vice President, Head of AI Strategy and Solutions

A Framework to Evaluate AI Maturity

A systematic approach to unlocking organizational maturity for AI

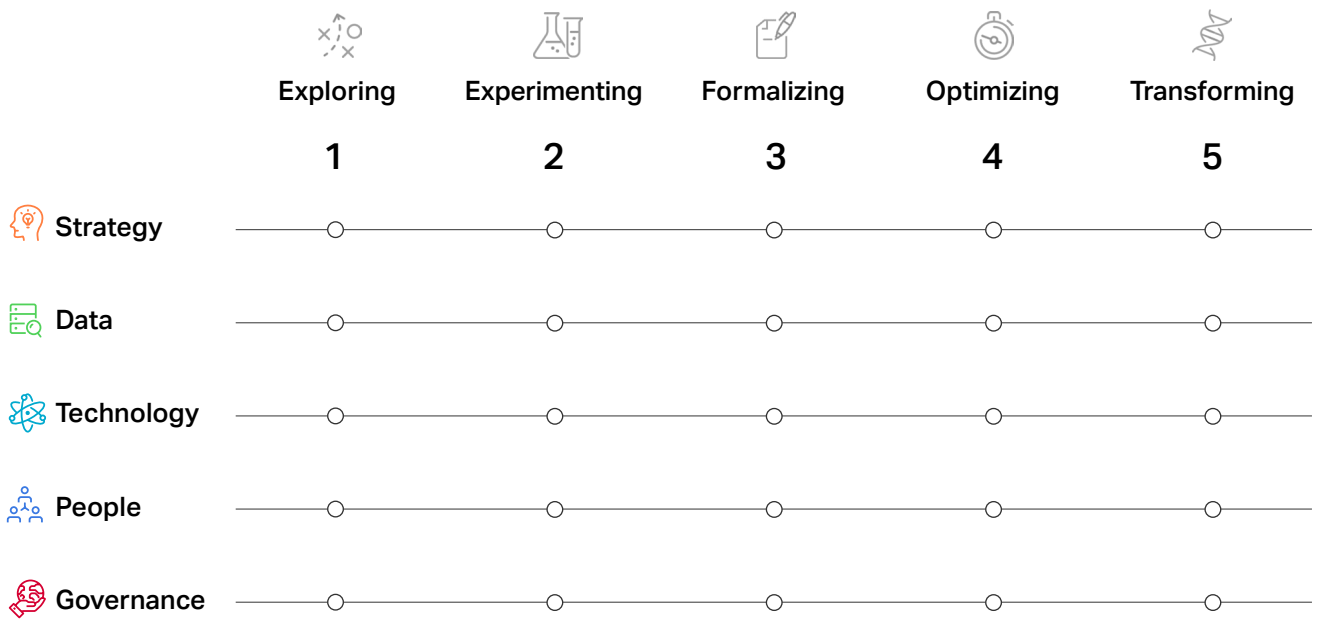

The AI Maturity Framework

AI is complex and multi-faceted, and to be applied, requires multiple parts of an organization to operate interdependently. In researching the state of organizational AI maturity in the industry, we were able to identify the five dimensions that an organization needs to update for AI and how those dimensions work together to enable and scale impact from AI over time.

Once we identified the key dimensions that define organizational AI maturity, we realized that there were few resources to help understand them. So, we designed an easy to understand framework to help organizations assess their ability to adopt artificial intelligence and decide what to do next.

The framework is a 5×5 grid that shows the relationship between the organizational dimensions needed to make AI real and the five stages of maturity that organizations go through as they level up these dimensions. The five organizational dimensions of AI maturity are Strategy, Data, Technology, People, and Governance. Each dimension is integral. A lack of progress in one will hold back overall progress on AI, even if other dimensions are further along.

For example, take an organization that has invested in a data lake and GPU (Graphical Processing Unit) cluster for AI. They also have a skilled data science team. But they have not set a clear business case

for AI, nor have they evaluated factors for securing trust with potential users. In this case, even the most sophisticated AI solution would fail to create value.

The takeaway is that time spent on Data, Technology, and People in this example is not wasted—but a lack of progress in Strategy and Governance delays time to ROI. Too many organizations today are either failing to anticipate hurdles across all dimensions or are over-preparing individual dimensions. Both slow progress.

The five stages, on the other hand, are simply inflection points on an organization's journey to achieving impact with AI.

At first, Exploring organizations must spend time understanding what AI can really do and how it could be of value for them. Experimenting organizations find out what will actually work and at what cost. Formalizing organizations are putting their first models into production with clear performance metrics, and typically, they use this process to drive additional investments. Optimizing organizations are focused on building out their ability to select, deploy and manage running AI solutions in production. Finally, Transforming organizations are using AI to push the boundaries of the technology and their own strategy.

The best way to move forward, wherever you are today, is to do a scan of your organization to determine which stage you're at based on the state of each dimension. Then, you can determine which dimensions will provide the critical leverage you need to move forward. From there, it's straightforward to design projects and work plans that move you forward.

The State of Organizational AI Maturity

In 2019, multiple studies showed that organizations were struggling to realize their vision for AI. In July, for instance, MIT Sloan Management Review found only 7% of organizations had put an AI model into production. Our own observations echoed these findings, so we took steps to learn more.

First, we created a survey to help organizations rapidly self-assess their organizational AI maturity across the five dimensions. Then, we used the survey to gather a purposive sample of senior decision-makers at large organizations in multiple industries in the U.S. and Canada, to create an up-to-date snapshot of AI maturity in industry.

Figure 1: Distribution of organizations by stage of AI maturity

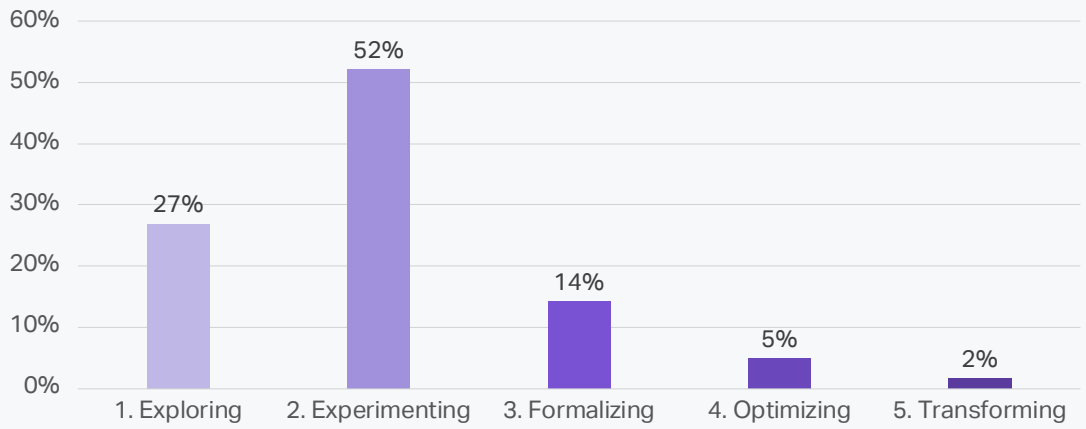

As shown in Figure 1, over a quarter (27%) are still trying to understand what AI means for their organization in the Exploring stage. A slim majority (52%) are in the Experimenting stage with AI and are working, either independently or with outside services or vendors, on AI Proofs of Concept (POCs). Another 14% are actively focused on putting a chosen AI solution into production in the Formalizing stage. Just 7% are at a level where they can reliably put solutions into production at scale.

Further insights are presented in the following sections and the survey is freely accessible for anyone to quickly snapshot their organization's AI maturity:

TAKE THE SURVEY

The Five Stages of AI Maturity

The stages of achieving business impact with AI solutions

An organization's stage of AI maturity determines the business value it can unlock from AI solutions. Although this stage is determined by the combined progress of five organizational dimensions, each stage shares similar challenges and opportunities that cut across dimensions.

Understanding the five stages helps you put your organization's current AI capabilities in context, including what your capabilities can help you achieve now (and what they can't) as well as what to anticipate for how those capabilities should develop in the future.

The five stages are:

STAGE 1

Exploring

Exploring what AI is and what it can bring to your organization. The organization does not yet have an AI model or solution in production.

STAGE 2

Experimenting

Experimenting with Proofs of Concept (POCs) and pilots. The organization is trying to put AI into production and can do so in limited ways.

STAGE 3

Formalizing

Moving from POC/ pilot to an AI solution in production. Putting AI solutions into production still requires significant organizational work at this stage.

STAGE 4

Optimizing

Scaling AI solution deployments efficiently as the number of deployed AI models increases. The organization is approaching a factory of model production.

STAGE 5

Transforming

Transforming the organization itself through the use of AI. The organization uses AI in how it operates across many critical areas of the business.

Figure 2: Distribution showing percentage of organizations in each industry that have reached each stage of maturity

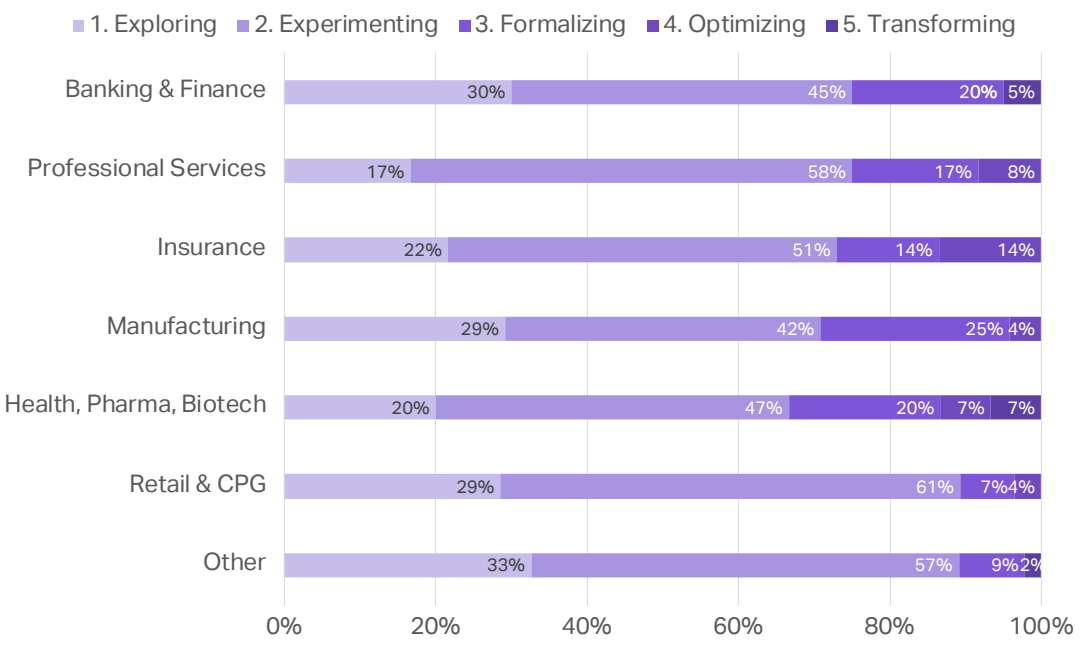

From our survey, we were able to gain insight into AI maturity stages across industries.

In our data, 90% of Retail & CPG organizations and 90% of organizations in "Other" industries are still Exploring and Experimenting. This number exceeds the baseline for all industries, and indicates that retailers are falling behind.

In Healthcare, Pharmaceuticals and Biotech as well as in Banking & Financial Services, 7% and 5% of organizations respectively had reached the Transforming stage. Banking and Financial Services also had the largest concentration of organizations that were still in the Exploring stage (30%), potentially signalling that progress is uneven in a way that may disadvantage Banking and Finance organizations still at this stage.

Manufacturing (25%), followed by Banking & Financial Services (20%), had the greatest concentration of organizations in the Formalizing stage. Organizations from these industries stand to gain the most from the AI Maturity Framework as they seek to align organizational dimensions to put AI in production.

STAGE 1

Exploring

Learn what AI can do and how to judge good AI opportunities from bad ones

In this stage, your organization is exploring what AI is and what it can bring to you. Your organization does not yet have an AI solution in production, but organizations with greater technical ability may start pursuing first Proofs of Concept (POCs) with AI.

Figure 3: Exploring stage by industry

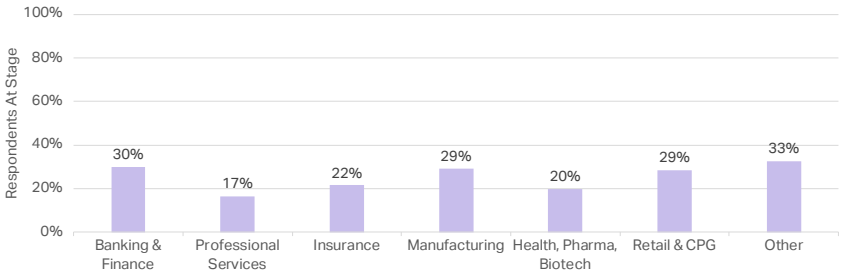

First

Organizations start Exploring when they make the shift from general awareness of AI to targeted questions about problems or opportunities that it can help them address. This might start with zero budget or with a formal charter for adopting AI. Either way, teams are still learning about specific benefits of AI for their industry and are unsure of how to realize them.

Next

Exploring tends to be driven by ambitious individuals or teams who focus on building informed interest and buy-in. They make progress by evaluating business use cases, costs, and benefits. Technical teams might start on AI experiments, but mostly as a tool for learning and creating internal awareness and excitement.

Later

Organizations reach a tipping point when they gain the ability to recognize good AI opportunities from bad ones. This allows teams to start building a roadmap of what work is required to define compelling AI solutions.

From analytics to AI at a financial institution

When a large financial institution went looking for potential applications of AI, it found hundreds. On closer analysis, dozens of use cases weren't true AI projects, but were addressable using traditional CRM solutions, business process automation, reporting and advanced analytics. Other cases were not aligned to the institution's strategy. Business and technical leaders validated the remaining cases for desirability, feasibility, and viability to identify the most strategic options. By working together, they also developed a shared vision for AI longer-term. The strategic roadmap that resulted from this work provided the clarity and budget needed to advance to the next stage of AI maturity: Experimenting.

STAGE 2

Experimenting

Placing calculated bets to determine which AI opportunities are ready for production

In this stage, your organization is experimenting with Proofs of Concept (POCs) and pilots. The goal of these efforts is no longer to experiment, but to drive measurable business impact. Successful experiments help teams to build momentum for AI and create limited business value along the way.

Figure 4: Experimenting stage by industry

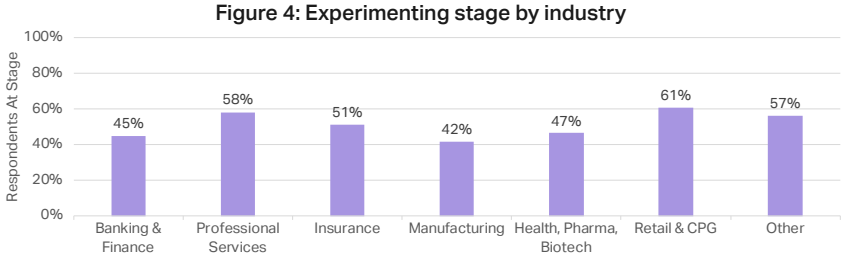

First

Organizations enter the Experimenting stage when they start testing hypotheses about what value can be created from specific AI solutions, and how. Usually, this is done with a Proof of Concept (POC). POCs might start with an AI software vendor or a single internal team able to operate independently.

Next

Experiments yield progress as their results clarify how to create business impact with AI out of the unique resources, opportunities, and challenges of the organization. This iterative learning approach is as much about verifying what AI can actually do as it is about clarifying what else is required to achieve impact. Teams that make the swiftest progress are careful to maintain focus on identifying blockers and enablers for AI models in production, especially AI governance topics like reliability, safety, trustworthiness, and accountability.

Later

Experiments might yield business value when deployed as a calculated risk into a limited application area. It's more important in the Experimenting stage for teams to develop a good handle on which projects should be put into production and how they will measure success.

Proving the case for straight-through insurance claim processing

At an insurance company, processing insurance claims at scale was a growing challenge. New, deep-learning-based Optical Character Recognition (OCR) techniques looked helpful for intaking claim forms faster. New predictive techniques looked beneficial for streamlining claim approval. Still, they needed to know what level of performance would be possible, at what cost, for their unique market niche. They curated a set of test data and performance metrics to carefully evaluate trade-offs such as rates of false negatives and positives. The experiment yielded a gradient boosting model that could safely increase straight-through processing rates and save up to 27% of current processing costs. This Proof of Concept (POC) allowed the insurer to design a pilot project in production for the next stage of maturity: Formalizing.

STAGE 3

Formalizing

Piloting AI solutions and aligning the organization to move ahead

In this stage, your organization is formalizing its efforts in AI by deploying pilot projects into production with a user adoption plan for achieving target performance metrics. The goal of these efforts is no longer to discover what AI could do in the environment, but to drive measurable business impact with it.

Figure 5: Formalizing stage by industry

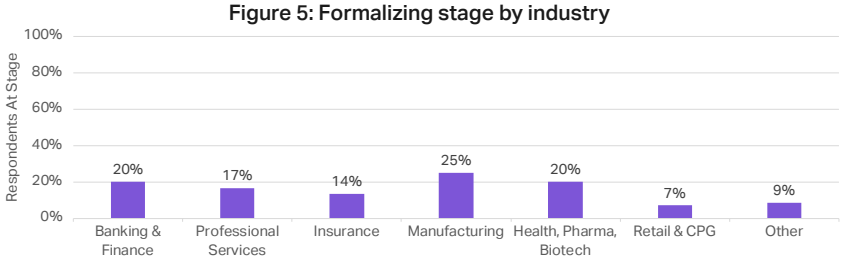

First

Organizations enter the Formalizing stage when they successfully deploy their first AI projects into production, usually as limited pilots. The goal is no longer to experiment to find what will work, but to leverage the lessons and outcomes of the experiment for measurable business impact.

Putting AI solutions into production requires significant effort at this stage, so each solution must have a clear business case with agreed-upon performance metrics. Additionally, internal risk policies and industry regulations simply won't allow AI projects to go live without adequate processes and relevant software tools to ensure their responsible use. If the organization has not yet matured in AI Governance, it quickly discovers gaps at this stage.

Next

Initial AI solutions might be budgeted, developed and deployed in an ad hoc manner to start, but Formalizing organizations use their experience to refine future plans for standardizing or streamlining AI delivery.

This focus guides the organization to confront any dimensions that it has not yet developed. For example, the data required to run an AI solution in production might necessitate expensive, bespoke system integrations, raising awareness about the need for more integrated data strategy.

Later

To adopt more complex applications of AI in critical business processes, executive-level sponsorship helps to increase budgets, mandates and plans, with special attention paid to ensuring AI models are safe, responsible and maintainable over time.

Formalizing a machine learning model to reduce delays in the transportation and logistics industry

Trucks load and unload cargo ships at a rate of multiple times each per day. However, scheduling was a growing challenge, with many trucks spending hours waiting in queue on most days. Traditional statistical methods had uncovered some key factors causing delays, and a Proof of Concept showed that an AI model could use these features to double the accuracy of predicted wait times.

For AI to actually play a role in minimizing wait times for drivers, a finished solution would have to take into account the diverse needs of truckers, workers, planners, and transportation operating systems. First, a data audit confirmed the availability and quality of data for training and deploying machine learning models. This process also helped clarify requirements for technical system integrations. In-person interviews clarified how the problem was experienced by different parties, and at the same time, built buy-in for solving the problem with AI. A machine learning model was then trained to predict the behavior of multiple agents and processes in order to visualize actionable insights for users.

Finally, the solution could start being piloted in a limited capacity. Alongside a production environment for the model to run in, a system was put in place to gather metrics on the quality and value of the predictive model throughout operational and seasonal changes. This and other factors ensured that the system could be steadily expanded as its benefits were proven and as stakeholders gained confidence in its use. With its first AI solution in place, the organization had the foundations it needed to scale impact at the next stage: Optimizing.

STAGE 4

Optimizing

Scaling AI and integrating it across the enterprise

In this stage, your organization is applying AI both in internal operations and in products, services, or other interactions with customers and suppliers. Multiple AI solutions are delivering business value with clear ROI. The organization can also move quickly from needs discovery to deploying in production. As a result, technical enablers and business processes are being put in place to safely govern AI at scale.

Figure 6: Optimizing stage by industry

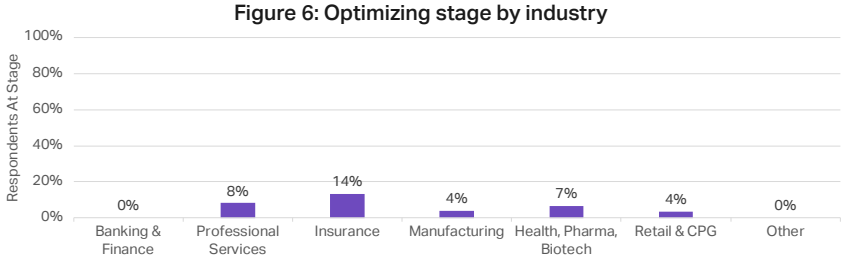

First

Organizations start to enter the Optimizing stage when they have at least one AI solution production and can reliably select, deliver, and manage additional AI projects with positive ROI.

Next

As the number of deployed AI solutions increases, new opportunities arise to improve the efficiency of delivering AI projects. For example, reusability of AI solution components and alignment between different organizational roadmaps allows for greater cost savings and faster deployment.

At the same time, new challenges arise around the complexity of supporting AI models in production. New infrastructure and programs are needed to integrate data, train users and to measure and control AI model performance at scale.

Later

The organization has completed investments to streamline the development and management of AI systems and has formalized policies and guidelines for using AI responsibly. Typically, C-level sponsorship has been involved to help drive integration across the organization.

Organizing for agile AI development at an insurance company

An insurance company had successfully deployed multiple AI models to production and wanted to scale their success across more of the business. The biggest blocker they identified was data preparation. Plenty of data was available, but data scientists and engineers were spending significant amounts of time organizing and analyzing data over the lifecycle of their AI solutions. To move forward, interviews and workshops were conducted to clarify the problem and define a shared solution designed to scale with future needs. They identified a lack of standardized methods and documentation for data analysis to be a key bottleneck, and mapped out new tools, processes, and technical as well as non-technical roles to address this challenge. Their new strategy to enable people, data and governance for AI helped the insurer close skill gaps for teams and prioritize investment in a data lake platform for streamlining AI model development. Every step of their plan to streamline AI delivery takes the insurer closer to the next stage: **Transforming**.

STAGE 5

Transforming

Actively shaping the organization with AI in new and profound ways

In this stage, your organization is pushing the boundaries of your industry and producing state-of-the-art work using AI. Your organization is not only applying AI to automate and augment business processes, but also to bring new business models, products or services to market. The organization has broken down organizational silos to integrate data and reimagine how value is created. AI drives decisions across the organization, supported by interconnected systems that learn and adapt over time.

Figure 7: Transforming stage by industry

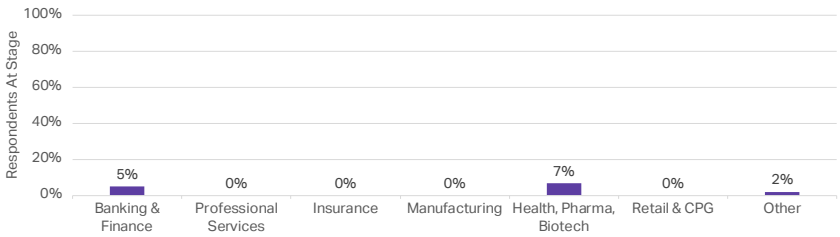

First

Organizations enter the Transforming stage when all organizational enablers are in place for AI and the majority of business decisions can be made with or by artificial intelligence. Widespread AI literacy and successful communication of the AI vision and roadmap have enabled support for working across teams and breaking down silos to build next-generation AI solutions.

Next

The organization is using AI to actively define or redefine business models, products and services, in addition to operations. AI is a key budget priority. Executives base the majority of their decisions on AI-driven insights and the strategic direction of the company is closely linked with its use of AI. Organizational silos are breaking down further integrate data, infrastructure, talent and operations for AI.

Later

Once transformative AI maturity is fully realized, the technology is pervasive in business operations and across whole value chains, making it fundamental to how new strategic opportunities are ideated and implemented. Organizations that want to continue transforming must continue to advance the science and engineering of artificial intelligence as well as its ethical use in society.

Many paths to transformative AI

Few organizations in the world today have reached the Transforming stage and it's unclear if any are yet delivering on the full potential of this stage. Across industries, they tend to either be built around AI from the start or (re)built around digital operations before making a strategic shift to AI. The first category of "AI-first" firms includes platforms Uber and Airbnb as well as new R&D firms in advanced industries like aerospace and biotech. The second category of "AI-focused" firms includes giants Google and Amazon, which were digital-first since their dot-com inception, as well as incumbents like Microsoft, which had to invest heavily in digital transformation before transforming with AI. Today, most large organizations are in the second category and still need to make a significant shift to digital-first operations before unlocking transformative AI.

The Five Dimensions of Enterprise AI

Levers to upgrade AI organizational maturity

Organizations must change how they think, act and learn in order to take advantage of AI. The five dimensions represent the key areas of any organization where management practices, operations and infrastructure need to evolve to realize this change.

To successfully increase an organization's overall stage of maturity for AI, each of these dimensions must mature individually and together. The weakest link limits overall progress. By improving capabilities in less mature dimensions, business leaders can unblock progress for AI projects as well as accelerate their overall organizational maturity.

The five dimensions are:

DIMENSION 1	DIMENSION 2	DIMENSION 3	DIMENSION 4	DIMENSION 5
Strategy	**Data**	**Technology**	**People**	**Governance**
The plan of action for achieving the desired level of AI maturity in the organization.	The data required to support specific AI techniques defined by the AI strategy.	The technical infrastructure and tools needed to train, deliver and manage AI models across their lifecycle.	The leadership practices as well as roles, skills and performance measures required for people to successfully build and/or work with AI.	The policies, processes and relevant technology components required to ensure safe, reliable, accountable and trustworthy AI solutions.

In the AI Maturity Framework survey, we designed questions to measure organizations' progress in each dimension individually. For example, if a respondent indicates (1) "some teams or business units (BUs) have an initial AI strategy supported by their business leader" and (2) "we just started to train and develop AI models through Proofs of Concept (POCs)," their resulting score would be 50% of the total score possible in the Strategy dimension.

Using this technique to score organizations in each dimension, we were able to gain insight into AI maturity dimensions across the five organizational AI maturity stages, to see how they evolved over time.

Figure 8: Dimensional maturity score as a portion of total possible dimensional score, subset by stage of maturity

From this boxplot (Figure 8), we can see that organizations are typically driving progress by investing heavily in Strategy for AI. The momentum created by Strategy helps drive progress in other dimensions at each stage.

For example, 50% of organizations in the Exploring stage scored between 20-30%, indicating that, at most, some teams have an initial strategy and/or first AI use cases have been identified. In contrast, only a quarter of Formalizing organizations are still clarifying use cases or setting an initial strategy. Most at this stage are already considering enterprise-wide AI strategies. For organizations that want to accelerate progress, prioritizing the Strategy dimension can help clarify work in other areas at each stage.

Survey data also shows that Governance remains underdeveloped across most stages. Technology remains relatively immature in Exploring and Experimenting until a leap forward occurs in Formalizing.

Insights like these, which are explored in more detail throughout the following sections, present opportunities for leaders to look ahead at what roadblocks will stall progress in future—and pinpoint a plan of action to move forward with less friction by leveling up sooner.

DIMENSION 1

Strategy

Organizational vision and roadmap to sustain forward momentum for AI

Strategy at its core is about the choices that a business makes to win. Strategy for AI maturity focuses on the plan of action designed to achieve the desired level of AI maturity in your organization.

Your plan needs to offer clarity about what needs to happen to implement AI, where, when and why—including how the organization intends to win with AI once implemented. The choices required to make this plan need to balance short- and long-term goals, taking into account the current stage of AI maturity, the competitive landscape, the business's strategy and ambitions, and leadership's desired velocity for progress.

When organizations overlook the Strategy dimension, AI experiments lack the business direction and justification to overcome hurdles to deploying in production or staying relevant to the business after deployment.

Strategy for AI adoption, which is the focus of the AI Maturity Framework Strategy dimension, is not the same as organizational strategy using AI. However, the two are linked in that the long-term vision for how the organization will work and compete in the future using AI should inform plans for where to focus AI efforts. The challenge for leaders is that the ability to envision meaningful strategic moves with AI requires some AI literacy to start with. To get started understanding how to judge a good AI opportunity from a bad one, see our article, Why you need intelligent AI adoption.

Figure 9: How important is artificial intelligence (AI) to your organization currently?

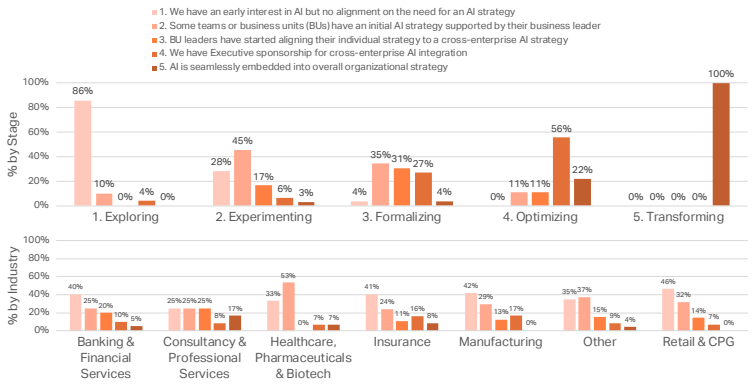

Figure 10: How prevalent is AI in your organization currently?

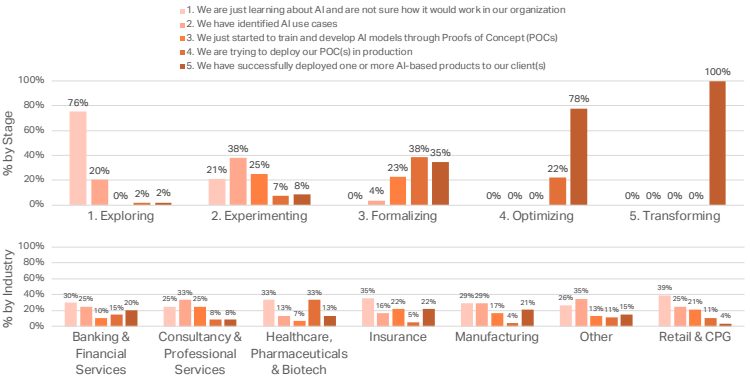

To develop your strategy for AI, consider:

1. **AI Maturity:** What is your current level of maturity for AI and any distinctive strengths or capabilities for AI in data, technology, people and governance?
2. **AI Trends:** What possible future scenarios involving AI could create disruptive benefits or challenges for the business?
3. **Horizontal and Vertical Alignment:** Do decision-makers at all levels have a shared understanding of AI and a shared vision of what opportunities to pursue with it?

The following sections describe the Strategy dimension at each stage of maturity and what organizations can focus on to level-up.

1-1 Exploring

Strategic alignment does not yet exist for what the organization wants to achieve with AI or how to achieve it. Usually, internal experts or enthusiasts are studying use cases or experimenting with personal side projects. These early visions for AI tend to be either too narrow (focused on non-critical parts of the business) or too broad and unrealistic, leaving projects without the value proposition or resources to proceed.

To move forward to Experimenting:

- Align business and technical leaders on the need for AI strategy to move forward

1-2 Experimenting

Organizations still have not aligned on an overarching strategy or vision for AI, but they are starting to do so in two ways: first, by planning how to use AI in a subset of the organization, such as a business unit or team; second, by refining and testing hypotheses about what business problems AI could solve using trials and Proofs of Concept (POCs). Typically, some executive sponsorship exists to unlock budget for POCs but the burden is on project owners to prove opportunities are worth investment.

To move forward to Formalizing:

- Align and galvanize leadership on AI investments using successful Proofs of Concept (POCs)

Refining AI concepts and developing a user research and data audit plan

A manufacturing company created a shortlist of interesting AI use cases to explore. Each use case looked plausible in the context of industry trends, and they couldn't decide on how to prioritize the next steps. To align plans for a Proof of Concept, they first needed to understand the technical feasibility of each use case in more detail. A concept framing document, user research plan, and data audit plan were developed to assess each use case in turn. This helped teams align on which use cases were most feasible, desirable and viable so they could move forward.

1-3 Formalizing

Executive sponsorship helps business teams to define the AI strategy for their organization. This sponsorship usually comes from a VP-level executive or above. With a clear strategy in place, little return may yet be realized from AI investments but the organization is able to make clear projections of ROI into the future. This enables the organization to unlock the budgets and mandate to execute their strategy.

To move forward to Optimizing:

- Document AI strategy for the organization to secure shared understanding
- Gain budget and C-suite sponsorship for AI projects

1-4 Optimizing

The organization starts executing against a clear AI strategy and mandate. C-level sponsorship usually exists to integrate AI across the broader enterprise. The AI roadmap is getting aligned with digital transformation, innovation, research and development, human resources, and other strategies. As a result, budgets are pre-approved and earmarked for AI initiatives across most business units. ROI for AI solutions is measured with formal metrics and recognized in the fiscal planning process.

To move forward to Transforming:

- Align AI strategy with other organizational roadmaps
- Discover opportunities to coordinate AI efforts across functions for greater impact

1-5 Transforming

AI is seamlessly embedded into overall organizational strategy. Budget schemes and indicators for both business and AI technology are integrated, enabling organizations to readily discover and act on new AI-based operational improvements and business models. The organization has the experience required to envision major innovation to their work, products, and services over longer time horizons of multiple years.

To keep moving forward:

- Sustain momentum to keep innovating and transforming

DIMENSION 2

Data

Powering AI models from training to production

Data for AI maturity refers to the enablement of usable data required to train AI models. No data, no AI. But how much is enough? In fact, different AI techniques require different types and amounts of data. Simulation-based modeling does not require vast sums to start, and synthetic data can be used to augment smaller datasets. Therefore, opportunities found in the organization's data should influence the design of the AI roadmap, but formal data requirements should be defined by AI solution requirements, not the other way around.

Typically, data requirements include considerations for being clean, complete, labeled (if using supervised machine learning techniques), integrated, secured, and corrected of harmful bias. These requirements apply across the full lifecycle of AI development, from training and testing to maintenance and retraining in production. They also enlist both technical and business ownership to successfully manage. For example, business users of systems that generate data must understand the downstream effects of using systems differently over time.

Today, the main challenge for most organizations is not a lack of data, but a lack of accessible and useful data for the AI solutions they wish to implement. To move through the Experimenting stage, most (69%) are struggling to gather data or are collecting and cleaning it in an ad hoc manner (Figure 11). In the Formalizing stage, a smaller majority (58%) have learned from this experience to build dedicated practices and infrastructure to support multiple AI solutions (Figure 11). For example, one in five Formalizing organizations have a standard data cleansing and consolidation pipeline (Figure 12).

Data is not destiny when it comes to AI, but it does have a role to play in drawing the AI roadmap. Data that's clean and readily accessible can help teams experiment faster. So-called "low-hanging fruit" generates early momentum and buy-in as the organization learns more about its data and about other dimensions for AI. To move beyond easy wins, you need to find or create data that captures unique elements of the business. When unique organizational know-how is captured in data, AI solutions can more easily create distinctive value for the organization and reinforce competitive advantage.

Figure 11: Are you able to access all the data you need for AI?

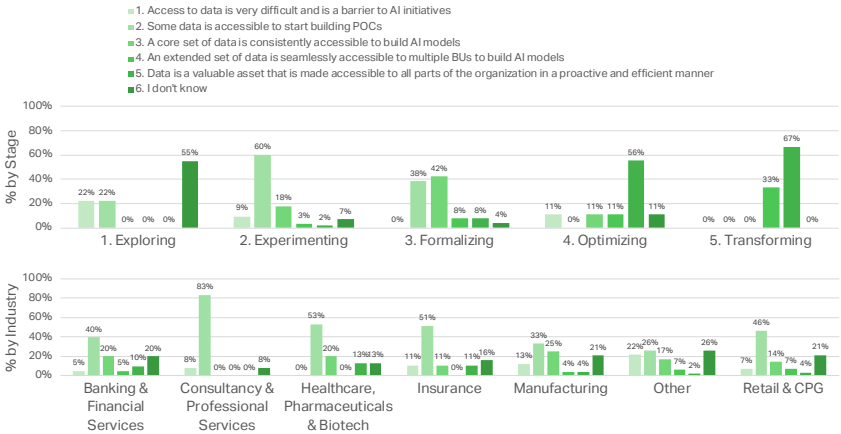

Figure 12: Is accessible data cleaned and consolidated for use with AI?

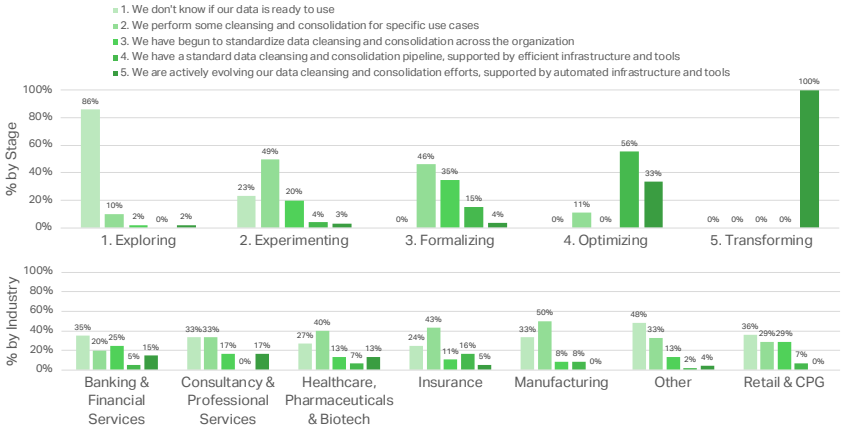

To prepare data for AI at an organizational level, consider:

1. **Volume:** Is there sufficient data to support AI techniques suggested by the AI roadmap?
2. **Representativeness:** Is there sufficient data to capture the range of situations that will be encountered by the use cases found in the roadmap?
3. **Quality:** Is data well-structured and free of gaps and errors?
4. **Labelling:** If using supervised learning techniques, is data labelled properly to enable AI models to understand examples?
5. **Accessibility:** Is data accessible for development as well as production environments?

The following sections describe the Data dimension at each stage of maturity and what organizations can focus on to level-up.

2-1 Exploring

Three main challenges impact the use of data for AI: visibility of internal datasets is low, special expertise is often required to understand data once it is found, and no standard infrastructure or process is in place to ease access to data. For example, structured data is often stored transactionally in databases and records, and it is siloed across the organization's different departments. Furthermore, the organization is not able to define data requirements for AI effectively and does not have clear plans to consolidate

data. For example, leaders do not have a good sense of what unstructured data sources could be available for AI.

To move forward to Experimenting:

- Learn about the data requirements for different AI techniques
- Look for unique elements of the organization captured in data to help inform the strategic AI roadmap

2-2 Experimenting

By learning more about data requirements for AI, teams have been able to assemble some data in a usable and accessible format. Some efforts may be underway to create common data stores or data lakes, but typically only limited data sources are connected, data is only refreshed periodically, and users have limited access. Specialised tools for data preparation, such as for data labelling, make a more immediate impact on preparing data for AI models.

To move forward to Formalizing:

- Use first AI experiments to build support for breaking down data silos and consolidating data

2-3 Formalizing

The organization has a core set of usable data that is accessible to build AI solutions. This success owes less to a generic strategy of gathering all data across the organization than to targeted, prioritized data collection based on a strategic roadmap of AI use cases. However, data enablement has been identified as a strategic priority, unlocking budget for building or growing common infrastructure (such as a data store or data lake) or for obtaining new data (such as by labelling existing data or installing new sensors for data capture). The organization can reliably measure the quality of data for specific AI techniques and use cases.

To move forward to Optimizing:

- Continue to break down data silos with AI use cases in mind
- Define metrics, processes and technologies for managing data quality for AI

2-4 Optimizing

Organizations have extensive, up-to-date, usable data to build complex AI solutions across the business. A majority of strategic systems are connected to a common data platform and are actively synchronizing information to the platform and between each other. The data platform is widely socialized within the company and accessible using intuitive graphical interfaces. Visibility and expertise on all internal datasets is

Data strategy proceeds from organizational strategy

An insurance company wanted to scale up their AI capability in order to make their products more valuable to customers in an increasingly digital world. They were in the midst of undertaking a major initiative to overhaul IT infrastructure and double the size of their data science team. They worked across multiple departments to define an AI Data Strategy that would lay the foundation required to embrace disparate data sources in alignment with best practices, ensure scalability, accelerate implementation and maximize the value of AI within the business. This collaborative approach created buy-in and alignment between business and technical stakeholders to drive rapid progress, with some recommendations implemented only weeks after the strategy was defined.

significant, and streaming data pipelines allow real-time access for priority use cases. The organization starts to actively clean and prepare data based on quality metrics aligned
to the AI roadmap.

To move forward to Transforming:

- Further automate, aggregate and make accessible data as efficiently as possible
- Identify new technologies, processes or partnerships needed to acquire new data

2-5 Transforming

The data platform is fundamental to how the core functions of the business operate, therefore, the infrastructure and tools to consolidate data are highly automated and empower teams to easily ingest new datasets. Data is well documented and both internal and external datasets have high visibility. Strategic investment ensures a self-service process for accessing data, from data ingestion to data consumption. Health monitoring of the central data repository is highly automated and provides real-time, reliable monitoring with minimal human intervention.

To keep moving forward:

- Get the most out of existing data with new AI techniques
- Continue to look beyond existing systems for new sources of actionable data

DIMENSION 3

Technology

Tools, infrastructure and workflows for powering AI across the solution lifecycle

Technology for AI maturity refers to the tools, infrastructure and workflows required to support the entire AI solution lifecycle, from training and testing, to deploying and running in production, to monitoring and retraining over time. All AI solutions share this lifecycle, whether purchased or built by internal teams. Leaders need to understand how technology is supporting each step in this lifecycle and what trade-offs are being made along the way as the organization matures. For example, a server environment that supports one AI model in production may not scale at a reasonable cost to multiple AI models.

For most organizations, the two biggest areas of technology change are development tools and computing hardware. New development tools include AI frameworks like TensorFlow and PyTorch. They also include software categorized with terms like DevOps, MLOps, and AIOps. These follow a broader industry trend of enabling closer collaboration between engineering and infrastructure management practices; the need for iterative development in AI model training accelerates this trend. New computing infrastructure, including purpose-built AI chips or GPUs (Graphical Processing Units), leverage chip architectures that are better suited to AI algorithms than traditional processors.

Today, while it is increasingly easy to start AI experiments on personal computers, 45% of organizations in the Experimenting stage already have dedicated servers for AI solutions, with some starting to use AI to predict variations in server workload in order to scale resources automatically (Figure 13). Still, at the Formalizing stage, only about a third (35%) are monitoring AI models for governance issues such as concept drift, and only 13% have procedures for retraining and updating AI models in production, whereas both of these metrics jump considerably at higher stages (Figure 14).

In contrast to rule-based software configured with step-by-step instructions, modern AI solutions are configured by setting goals or objectives that shape processes of machine learning. This is why AI solutions need to be iteratively trained and tested during development as well as monitored and retrained in production. As the business environment changes over time, machine learning models can degrade in performance if they are not retrained. This problem can be combated by designing models to adapt continuously to new data, but at the expense of added complexity for AI governance, including techniques for monitoring AI models running in production.

Figure 13: Which of the following are true about the technical foundation for AI?

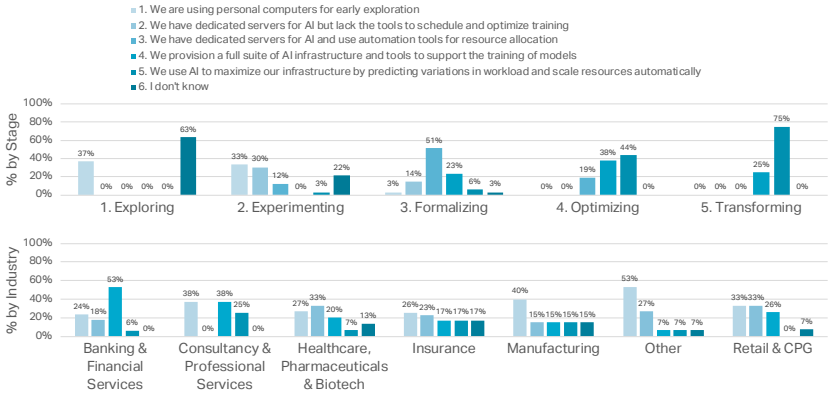

Figure 14: How are deployed AI models being maintained in your organization?

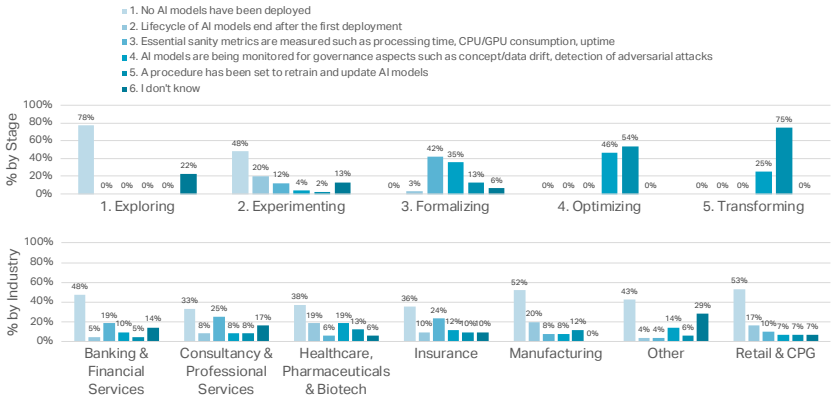

When building up technology for AI, consider:

1. **Requirements:** What is needed today, and how fast will those needs evolve?
2. **Flexibility:** How can tools connect to different types of data, support different types of modeling approaches and AI frameworks?
3. **Scale:** Can the technology scale to different production scenarios?
4. **Policies:** What policies are needed for this technology to function and succeed?

The following sections describe the Technology dimension at each stage of maturity and what organizations can focus on to level-up.

3-1 Exploring

Organizations typically don't have specialized AI or machine learning solutions in place, even when investments in adjacent technologies like DevOps, robotic process automation (RPA), or advanced analytics already exist. Business leaders are unsure of what's needed. Any initial experiments are conducted on personal computers or cloud-based environments.

To move forward to Experimenting:

- Determine what technology you need to conduct first AI experiments, starting with personal computers and cloud development environments

3-2 Experimenting

Data scientists and developers start using cloud infrastructure to share know-how and results, and to leverage GPU power beyond the confines of their laptops. Cloud-based or on-premise servers can be provisioned. AI model training happens manually with no automated resource management facility. If the organization has a DevOps team, they are likely not yet used to

deploying AI models. There is no standard process or deployment architecture.

To move forward to Formalizing:

- Formalize deployment architectures and look for ways to automate their use

3-3 Formalizing

In order to fully enter production, technical controls exist to allow any "human in the loop" and explainability features defined by AI governance practices. AI deployment architecture and development tools are standardized and implemented. Access and resource allocation for computing power is managed by an automated system. As the process of developing and deploying AI becomes more standardized and scalable, departments experiment with more complex AI solution designs. For example, the organization has an approach to reusing an AI model trained in one part of the business for a similar task in a different part of the business.

To move forward to Optimizing:

- Continue to streamline development tools and computing resource management

3-4 Optimizing

As the number of deployed AI models increases, organizations invest in new infrastructure to manage AI development, deployment, and management more efficiently. Management of deployed models includes retraining
on new data. Solving these challenges involves centralizing tasks such as monitoring and auditing AI models for compliance, performance management, or troubleshooting purposes, or supporting reuse of models and other code.

To move forward to Transforming:

- Invest in a centralized platform to track, deploy, and retrain AI models

3-5 Transforming

AI deployment architecture is standardized and efficient. As AI becomes more central to the organization's overall strategy, new use cases drive the organization to push the boundaries of technological capabilities to build state of the art AI solutions; for example, scaling to new locations might require specialized edge hardware, or personalization of AI models for individual customers and suppliers might require automatic provisioning of computing environments. These needs push the organization to use AI to manage the technology infrastructure itself. For example, the AI computing environment optimizes resource provisioning automatically.

To keep moving forward:

- Define innovative new use cases that push the boundaries of existing technology

Using AI to improve AI resource management

Specialized cloud infrastructure for AI was an indispensable early investment at Element AI, but we never guessed that our GPU (Graphical Processing Unit) processing needs would grow 36X in just 2.5 years. Recently, we surpassed 6 million total training jobs on the cluster. To operate at this scale, we started over a year ago to create an internal product for managing AI training jobs and optimizing use of our GPUs. By building our own solution, we could combine the distinct needs of our IT infrastructure team and AI scientists alike, from technical features like container orchestration and automatic scheduling, to non-technical design goals like ensuring fair access for scientists and low management overhead for IT. The solution now enables our 100+ fundamental and applied research scientists to run 2000-5000 training jobs daily. Meanwhile, our IT team manages 1400 GPUs with minimal increase in staff, while enabling less than 1% of AI jobs to be queued due to lack of resources.

DIMENSION 4

People

Roles, skills and measures of success to work smarter with AI

The People dimension of AI maturity focuses on aligning leadership and change management to ensure people are ready, willing and able to use AI. Even the most intelligent AI solutions will not succeed if people are not organized and motivated to use them. And it is the responsibility of executive leaders to help business and technical teams deliver AI and work successfully with it.

To lead people for AI, leaders need to help teams to bridge their expertise so they can generate the best vision, roadmap and day-to-day tactical decisions for AI. This means helping people at all levels to make a series of mindset shifts: from building rule-based systems with known development processes, to adopting learning systems that require iterative development and continuous care over time; and from doing work, to collaborating with AI systems that participate in work. Accordingly, leaders themselves must have a good grasp of the implications of AI for their business, so they can provide the best direction at all times and roll up their sleeves to make decisions when needed.

For people to successfully build and work with AI solutions, they also need training, job support and meaningful inclusion in the process of designing and deploying AI solutions. Training should address business as well as technical aspects of AI so employees can understand and help shape the organization's unique AI vision. Job support includes frequent communication of the AI roadmap and help with job upskilling and reskilling when appropriate. Involving users in the design and deployment of AI solutions builds trust and ensures that solutions leverage the best information available at every decision step. Across the whole organization, this is more about separating myths from realities about AI than about upskilling all employees to understand AI algorithms.

Our survey data suggests that as organizations mature, organizational support for People reaches a tipping point once multiple AI solutions are in production where informal practices are no longer adequate. Almost half (49%) of Formalizing organizations use organic "communities of influence" to upskill and drive adoption, but this number drops to almost nothing at higher stages. Meanwhile, just a third of Formalizing (33%) organizations, and almost no organizations at earlier stages, have invested in change management guidelines—whereas most Optimizing and Transforming organizations (88% and 75%) have done so.

A recurring question for business leaders is how AI will impact jobs. In fact, no one knows what will happen, but research supports the outlook that the choices made by leaders themselves have an integral role to play. For example, AI solutions can be applied for automation as well as human-machine collaboration, and even when used for automation, AI does not have to substitute for workers but can complement and even increase demand for their work. This means leaders can seek out uses for AI that better leverage the unique people, culture and values of their organization. For more about the impact of AI on jobs, see our article, Jobs and employment after AI: reasons for optimism.

Figure 15: Which of the following are true about the level of AI knowledge and expertise within your organization?

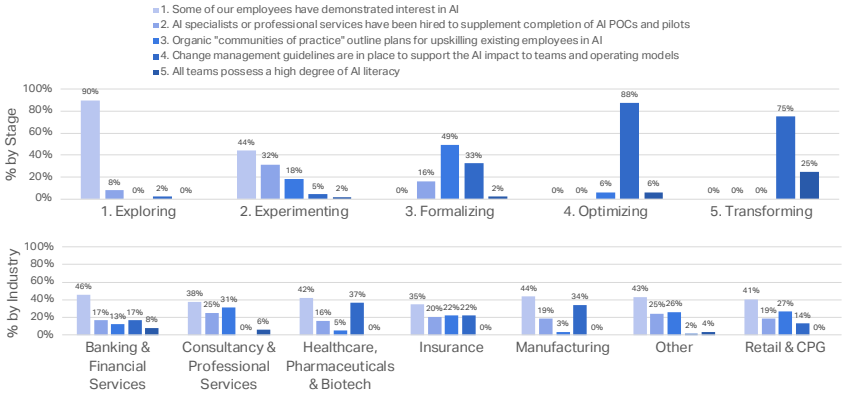

When preparing people for AI, consider:

1. **Leadership Persona:** Who is leading the effort to enable or scale AI? Is the right leader positioned and informed to make the decisions required?
2. **AI Literacy:** More than just technical training, do individuals have the ability to actively learn and adapt to AI technologies?
3. **Job Skills and Resources:** What people and parts of the organization need to reskill/upskill in order to adapt to changing demands and roles? What other resources will employees need once they start working with AI?
4. **Talent Strategy:** What new talent will be needed, and on what basis? What partners or other outside help should be used?
5. **Operating Model:** Who will manage AI resources, projects and solutions over the course of the AI roadmap? Should management of AI be centralized or decentralized?

The following sections describe the People dimension at each stage of maturity and what organizations can focus on to level-up.

4-1 Exploring

The organization hasn't defined roles and responsibilities for AI and doesn't yet know how to do so. In the short term, business teams need help absorbing the applicable takeaways from technical literature so they can build valid use cases for AI. Data science teams need help from business partners to connect AI techniques to a meaningful business problem or opportunity; they may also need help understanding AI techniques at a technical level.

To move forward to Experimenting:

- Develop AI literacy of business and technical teams to build confidence and support

- Encourage knowledge sharing between teams to ensure AI is accessible to all
- Enlist help from AI specialists to identify and address knowledge gaps faster

4-2 Experimenting

Some definitions exist for roles and responsibilities of individuals working with AI, but the organization is still experimenting to discover the right way to organize for AI. Typically, small teams with internal experts in data science, business intelligence (BI) or advanced analytics start experimenting with Proofs of Concept (POCs). However, organizations must resist the temptation to allow these teams to work in isolation. Instead, POCs should help the
organization discover what additional AI literacy (both technical and non-technical) is needed. For example, leaders should start communicating the AI vision and roadmap to employees, and people from different levels and functions should be enlisted to help define and conduct AI experiments.

To move forward to Formalizing:

- Assign cross-functional, flexible, networked teams to own AI experimentation
- Organize learning activities for AI, such as education, hackathons or secondments
- Identify AI career paths and implications for workforce planning activities

4-3 Formalizing

New roles for AI, such as machine learning engineer, have emerged and are being defined at the Enterprise level. Performance metrics are being established but are not yet used in formal performance management processes. Typically, organic Communities of Influence (CoIs) or a dedicated Center of Excellence (CoE) have been created to provide skills and resources for new roles, guidance on acquiring outside talent, and education for others in the organization. Business leaders are communicating the AI vision and helping to motivate and educate employees to share in that vision.

To move forward to Optimizing:

- Define AI accountabilities for executive leadership, team roles, structure and budgets to deliver against the AI roadmap
- Update rewards, recognition and performance standards in place to attract and retain AI talent
- Cultivate Communities of Influence (CoIs) or a Center of Excellence (CoE) to engage individuals outside the formal AI organization

4-4 Optimizing

Bridging business and technical teams to build useful AI products

At a financial institution, a team of machine learning engineers was hitting a roadblock: they had analyzed financial data and started experimenting with code, but they could not move forward in their attempts to work with financial analysts. After answering a few ad hoc questions about how they used data in their day-to-day work, analysts were refusing to answer more questions and lodging concerns with management to slow progress. Once leaders sought to understand the cause for this pushback, it turned out that analysts had assumed the worst about the intentions behind AI-based tools for their jobs. In fact, engineers were designing to help analysts with their toughest challenges—not automate their jobs. The project resumed in full force once leaders helped to address analysts' concerns and connect business and technical teams for better collaboration.

Organizations have clearly defined responsibilities and KPIs for new roles associated with AI. The broader talent strategy supports the learning journey of all employees to increase AI literacy and adapt to changes in work introduced by AI. The talent strategy includes plans to build specific AI capabilities and upskill or transition existing workforce as required. Leaders are actively involved in helping the organization adapt to change. Organizational structures like CoIsPs or a CoE are formalized and their mandates expanded to include managing the organization's relationship with the broader AI ecosystem, such as through vendor and partnership management.

To move forward to Transforming:

- Include representation of the AI organization at the executive table with accountability for Enterprise KPIs for AI
- Establish sustainable learning journeys for individuals responsible for delivering and using AI

4-5 Transforming

All teams and employees possess a high degree of AI literacy and promote a culture of working in complementary or collaborative relationships with AI systems. AI is integrated in some way for all roles, including at the executive level, and is likely to be used to help HR and talent teams to plan and operate. As a result, the organization's delivery model is transformed, changing how roles are defined and how people are expected to do their work.

To keep moving forward:

- Communicate self-driven career paths for AI to guide professional development in different areas of AI expertise
- Empower HR/Talent teams to use AI as a business transformation tool

DIMENSION 5

Governance

Policies, processes, and structures to ensure responsible and safe AI

Trust is the foundation of every interaction at your organization and AI governance is the foundation of trustworthy AI. Governance for AI maturity refers to the policies, processes and relevant technology components required to ensure safe, reliable, accountable and trustworthy AI solutions.

To deliver solutions that have all four of these qualities (as well as being high-performing), new forms of cooperation are needed between business, technical and risk teams. They need to connect practices from the design of AI solutions to the design of policies, process controls, and supporting technologies. For example, ensuring algorithmic decisions can be traced back to the data and models that produced them is important for debugging, compliance, and continuous improvement.

Today, organizations are less mature in AI governance than in any other dimension, and at the same time, the gap between the most and least mature organizations is widest. A large majority (72%) of all respondents either didn't know about their organization's governance efforts or indicated their company was just starting to learn about this important area. In contrast, 100% of organizations at an Optimizing level or above are invested in governance beyond what is required for mere regulatory requirement (Figure 16). Most (64%) had implemented metrics for tracking governance or even considered their AI ethics and governance practice to be a competitive differentiator.

Figure 16 What degree of governance has been established to enable AI?

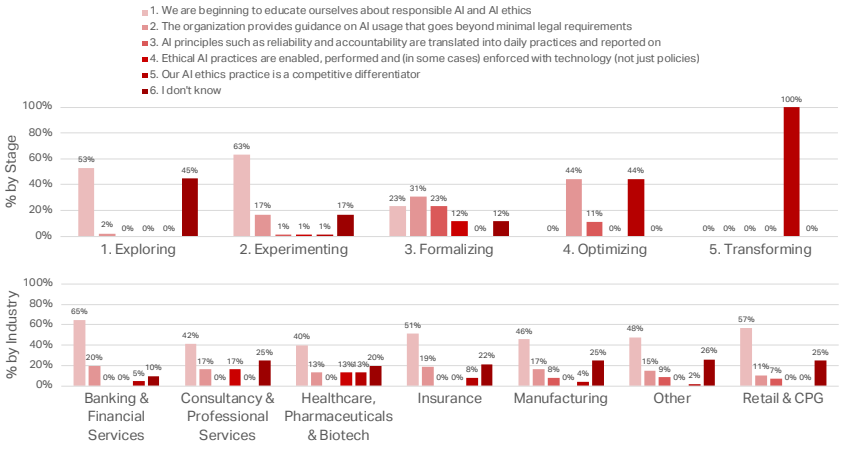

AI governance is about a lot more than risk management, but there are distinct new risks from AI systems that business leaders should know about. The unique risk profile of AI systems derives from the fact that, in contrast to rule-based systems configured with step-by-step instructions, AI systems are configured by setting goals or objectives that shape a process of machine learning. One way this process can fail is for machines to learn how to perform "test" cases using biased or incomplete knowledge that results in errors in the real world. Another is for AI models to be trained effectively to work in the real world, but to not be updated as the state of the world changes. For more about new risks from AI and how to manage them, see our article, How AI risk management is different and what to do about it.

When developing AI governance in your organization, consider:

1. **Risk:** In your set of use cases for AI, what are potential risks at each stage of maturity as the solution scales to more data, users, and impact?
2. **Regulation:** What relevant regulations do you need to pay attention to and follow in each country and jurisdiction that you operate in?
3. **Safety:** Does your AI solution protect or deal with people's personal safety? How can you ensure that your AI keeps people's safety as the highest priority?
4. **Explainability:** How can you demonstrate the reasoning that led to a prediction from an AI engine or model?

The following sections describe the Governance dimension at each stage of maturity and what organizations can focus on to level-up.

5-1 Exploring

Board members, management teams and employees are beginning to educate themselves about responsible AI so they understand new or heightened risks, obligations, and opportunities. When teams work together on a strategic roadmap, they can identify major risks associated with priority use cases.

To move forward to Experimenting:

- Understand new risks from AI such as model bias and model drift
- Identify specific risks along your AI roadmap and any new governance practices that might be needed above and beyond existing practices
- Start to develop high-level principles to guide responsible AI use moving forward

5-2 Experimenting

Business, technical and risk teams have a shared understanding of what's required for AI models to be compliant with any legal obligations related to AI across the solution lifecycle. The organization has begun to develop high-level principles to provide guidance on AI usage that goes beyond minimal legal requirements. To build trust, internal stakeholders who will use or be impacted by the AI system have a role in testing and refining the AI system design. To enable complex models to graduate to production, the organization starts exploring techniques like explainable AI (XAI) that help secure trust with users.

To move forward to Formalizing:

- Understand current debates about AI ethics and Fairness, Accountability and Transparency (FAccT)
- Involve different stakeholders to gain a complete view of potential challenges and opportunities for reliability, safety, trustworthiness, and accountability
- Translate principles into concrete role responsibilities, processes, and metrics

5-3 Formalizing

Guiding principles for AI governance are being translated into daily practices that track specific performance metrics for areas including safety, reliability, trustworthiness, and accountability. Reporting is centralized and key stakeholders have access to the data. Typically, a dedicated model evaluation function exists separately from AI modeling, similar to a QA team. The commitment to AI governance is formalized as a critical part of the overall AI strategy. Everyday practices that increase reliability and trustworthiness are part of the standard development cycle. External voices have been incorporated appropriately into discussions about AI ethics.

To move forward to Optimizing:

- Synthesize existing practices into guidance that generalizes to more use cases
- Investigate supporting technology for governance such as reporting tools

5-4 Optimizing

As the number of AI models deployed in production increases, so does the complexity of interactions between these models as well as the scrutiny from stakeholders and regulators on AI practices. To keep pace, responsible AI practices are guided by standard guidance and enforced through increasingly centralized and auditable processes, policies and technologies. The organization considers risk at the model and model portfolio level thanks to sophisticated understanding of the dependencies and feedback loops between people, different AI applications that are running in production and the business environment.

To move forward to Transforming:

- Build organizational structures to manage the strength and scalability of AI ethics and governance across multiple parts of the organization, such as an ethics board

5-5 Transforming

Strong governance has enabled the organization to go beyond regulatory compliance. Its multiple lines of risk defense and stakeholder trust are a competitive advantage for applying AI in powerful ways. This puts the organization in contact with novel challenges related to AI ethics or Fairness, Accountability and Transparency that may not have been confronted yet in their industry. The organization may formally invest in capabilities to drive multi-stakeholder agreement about how to navigate these challenges for shared benefit. This can include disseminating the technologies and approaches it has developed.

To keep moving forward:

- Engage with the broader AI ecosystem to help shape AI governance at the level of industry standards and best practices

Opening the AI "black box" to build trust and increase auditability

A manufacturer created a sophisticated AI model for anomaly detection. The manufacturer was committed to catching all defects before shipping to customers and the new model allowed plants to identify defects better than ever before. However, data science teams had implemented the model in a "black box" solution that could not demonstrate to plant quality inspectors why some parts were rejected while others were not. Without being able to understand the reasoning used by the model, inspectors resisted implementation. Data scientists had already proved the accuracy of their model to satisfy risk and compliance needs, but further investment in AI governance was required before stakeholders would find the system trustworthy and accountable. To move forward into production, AI explainability techniques are being used to add new features that visually explain the reasoning of the anomaly detection model.

Conclusion
Putting it All Together

Operationalizing AI is not simple. Many organizations either fail to anticipate hurdles in Strategy, Data, Technology, People or Governance—or, they over-prepare in a single dimension. Both errors slow down progress. Losing too much time can make the difference in successfully competing with AI over time.

But perseverance pays off. Using AI, leading organizations have already dramatically transformed, yielding incredible benefits for their bottom line, for society and for the future. The key is to start, one use case at a time, and stay the course until the organization can scale their operations and explore new products, services and business models for transformative impact.

And there's a systematic path to progress. In this document, you've learned about the Five Stages and Five Dimensions of AI Maturity. Across the five stages, from Exploring to Transforming, the dimensions act as levers to level up.

As demonstrated by multiple cases and by our survey results, this framework has helped our clients overcome specific challenges on their way to deploying and scaling AI.

1. A financial institution found the AI "needles" in a haystack of analytics and automation projects so they could start experimenting. **Strategy unblocked Experimenting.**
2. A manufacturing company used explainable AI (XAI) techniques to gain the trust and support of users to put an AI model into production. **Governance unblocked Formalizing.**
3. An insurance company defined a formal data strategy to scale AI delivery to new areas of their business. **Data unblocked Optimizing.**

In every case, progress in a less mature dimension allowed the organization's strengths in other dimensions to propel AI projects forward. Knowing which area needed work was half the battle.

The AI Maturity Framework will similarly help any organization, in any industry, to pinpoint which dimension is holding back their progress and how they can take action to move forward.

No matter your stage of maturity, therefore, don't hesitate. Check back on the framework as you progress. Remember, you are creating a living roadmap for adopting AI and realizing incredible benefits for your organization.

If you haven't yet, we encourage you to spend time with your team understanding current AI maturity in the five dimensions. Our free survey is available as a lightweight tool for doing this. Click here to take the survey:

TAKE THE SURVEY

ELEMENT AI

Executive Blueprint

	Exploring	Experimenting	Formalizing	Optimizing	Transforming
	1	2	3	4	5
Strategy Organizational vision and roadmap to sustain forward momentum for AI	Align business and technical leaders on the need for AI strategy	Use AI Proofs of Concept (POCs) to align leadership on AI investment	Define AI strategy and communicate to the organization Gain budget and C-suite sponsorship for AI projects	Align AI strategy with other organizational roadmaps Find ways to integrate AI across the organization	Sustain momentum to keep innovating and transforming
Data Digital sources of truth for training, running and improving AI	Learn about data requirements for different outcomes from AI solutions Identify unique elements of the organization captured in data	Use first AI experiments to build support for breaking down data silos and consolidating data Audit available data and data quality	Standardize data cleansing and consolidation across the organization	Further consolidate data with a focus on quality and efficiency Identify technologies, processes or partnerships needed to acquire new data	Get the most out of existing data with new AI techniques Continue to look beyond existing systems for new sources of actionable data
Technology Tools, infrastructure and workflows for powering AI across the solution lifecycle	Determine what technology you need to conduct first AI experiments	Formalize deployment architectures and look for ways to automate their use	Continue to streamline development tools and computing resources	Centralize AI model deployment, monitoring, resource management and retraining	Define innovative use cases that push the boundaries of existing technology
People Roles, skills and measures of success to work smarter with AI	Increase knowledge sharing between business and technical teams Enlist help from AI specialists to identify and close knowledge gaps faster	Sponsor cross-functional, networked teams to own AI experiments Connect AI teams with people at different business levels and functions	Cultivate AI Communities of Influence or a Center of Excellence Update talent strategy for new skills and roles required for AI	Incorporate technical and non-technical AI skills into your learning roadmap and talent strategy	Empower HR/ talent teams with AI solutions
Governance Policies, processes, and structures to ensure responsible and safe AI	Understand new risks from AI and current debates about responsible AI Identify specific risks along the AI roadmap Develop high-level principles to guide responsible AI use	Involve different stakeholders to gain a complete view of challenges and opportunities Translate principles into concrete responsibilities, processes and metrics	Synthesize existing practices into guidance that generalizes to more use cases Standardize governance practices with supporting technology such as reporting tools	Build organizational structures to manage the strength and scalability of AI ethics and governance across multiple parts of the organization, such as an ethics board	Engage with the broader AI ecosystem to help shape AI governance at the level of industry standards and best practices

GLOSSARY

AI Maturity: We use the term "maturity" to refer to the degree of formal, operationalized processes (in the context of AI) in place at the organization. In the context of software development, capability maturity models typically describe an organization's ability to deliver software projects in an optimized, repeatable way, especially through the use of controlling metrics. We argue that in the case of AI, organizational maturity shifts from merely delivering AI to delivering "with" AI at the Transforming stage of maturity.

Artificial Intelligence (AI): Artificial intelligence is human-like intelligence demonstrated by machines. Today, it is often used interchangeably with "machine learning" or "deep learning," but these are merely branches of research in a larger discipline combining aspects of computer science, neuroscience, philosophy, game theory and more. AI is popularly understood in terms of its applications like natural speech and language translation, autonomous vehicle operation, and simulation-based modeling. Practically, organizations benefit from AI by deploying it in specific solutions designed for specific use cases.

Explainable AI: Research area focused on developing AI models that are easier to understand and explain for human users. Alternatively, XAI can be defined as the process of extracting some form of explanation from pre-developed models that are otherwise difficult (if not impossible) to understand for their users.

Proof of Concept (POC): A project that helps the organization test a hypothesis about what business value is technologically possible to create and at what cost. It may run in a limited capacity in production but is usually limited to "test" data and systems. A POC implements a Use Case.

Use Case: A specific situation, task, and set of users/stakeholders and requirements that together define an opportunity to deploy AI for business impact. While many potential use cases exist in any business or industry, and some are more popularized than others (such as using a machine vision system for anomaly detection), a real use case is always defined for the unique context of a particular business.

We're excited for your journey to AI maturity and we're ready to be a partner on it.

Contact Element AI today if you'd like help accelerating your AI maturity today.

sales@elementai.com

Montreal
6650 Saint-Urbain,
Suite #500, Montreal,
Quebec, H2S 3G9, Canada
+1 (514) 379-3568

Toronto
296 Richmond St W
Suite #100, Toronto
Ontario, M5V 1X2, Canada

London
Vox Studios
Suite W-111
45 Durham, Vauxhall
London SE11 5JH

Seoul
Dreamplus Gangnam
311 Gangnam-daero
Seocho-gu, Seoul
Republic of Korea

Singapore
The Executive Center
10 Collyer Quay
Level 37 - Suite 11/15
049315, Singapore

ACKNOWLEDGEMENTS

This report was written by Karthik Ramakrishnan, Vice President, Head of AI Strategy and Solutions, and Cory Salveson, Advisory Research and Communications Strategist, with contributions from the following: Grace Abuhamad, Cyrielle Chantry, Simon-Pierre Diamond, Phil Donelson, Lisa Ebert, Walid Koleilat, Andrew Marble, Francois Ok, Martin Ryan, Charlotte Sobolewski, Patrick Suen, Sylvain Truong and Richard Zuroff.

Production team: Guillaume Gagnon, Jaxson Khan, Melissa Guay, Morgan Guegan, Simon Hudson and Wei-Wei Lin.