MIT Technology Review Insights

In association with

JPMorgan Chase & Co.

By focusing on data, value propositions, and talent, enterprises can get ready to take the next step beyond experimentation.

Successfully deploying machine learning

Preface

"Successfully deploying machine learning" is an MIT Technology Review Insights report developed in collaboration with JPMorgan Chase. The report is based on interviews with analysts, scientists, and senior executives of companies using machine learning and AI. The interviews were conducted in the autumn of 2022 to evaluate the integration of AI and ML into disparate parts of the enterprise, while scaling effectively and responsibly. Adam Green was the writer, Michelle Brosnahan was the editor, and Nicola Crepaldi was the publisher. The research is editorially independent, and the views expressed are those of MIT Technology Review Insights.

We would like to thank the following experts for providing their time and insights:

Eric Boyd, Corporate Vice President of AI Platform, Microsoft

David Castillo, Firmwide Head of AI/ML Technology and General Manager of AI/ML Product Line, JPMorgan Chase

Samik Chandarana, Chief Data and Analytics Officer for Corporate and Investment Banking, JPMorgan Chase

Drew Cukor, Firmwide Head of AI/ML Transformation and Engagement, JPMorgan Chase

Iya Khalil, former Global Head of AI Innovation Lab, Novartis

Aleksander Mądry, Director of MIT Center for Deployable Machine Learning and Cadence Design Systems Professor of Computing, MIT

Lidia Mangu, Managing Director, Head of Machine Learning Center of Excellence, JPMorgan Chase

Foreword

This is a pivotal moment for AI and machine learning. As a data scientist, I'm always thinking about this. But the rise in popularity of AI and machine learning in people's everyday lives, and the daily prominence in news reports today, highlights the notable innovations in AI and what they mean for individual consumers and some of the most impactful institutions in the world. Collaborating on this report with MIT Technology Review Insights was an impactful experience, as we had the opportunity to hear survey respondents and interviewees' perspectives on AI in business.

At JPMorgan Chase, we excel at solving practical problems using cutting-edge techniques. Our Machine Learning Center of Excellence (MLCOE) partners with business groups and their dedicated data science teams, combining their business knowledge with our deep AI and machine learning expertise to build and deploy high-impact solutions. To date, our MLCOE has focused on natural language processing (NLP) and speech recognition, as well as time series analysis, reinforcement learning, representation learning, recommendation systems, and most recently, knowledge graphs. All of these MLCOE skills support the many use cases at JPMorgan Chase that help us to create better, faster, smarter experiences for our clients and customers.

Of course, none of this happens without talent. I'm very proud to work at JPMorgan Chase, recognized as a financial services leader in AI and supporting more than 900 data scientists across our company.

I hope you find this report inspiring as your teams dive deeper into the capabilities available now, along with the new techniques that will emerge in this innovative field.

Dr. Lidia Mangu
Managing Director, Head of Machine Learning Center of Excellence, JPMorgan Chase

CONTENTS

01 Executive summary

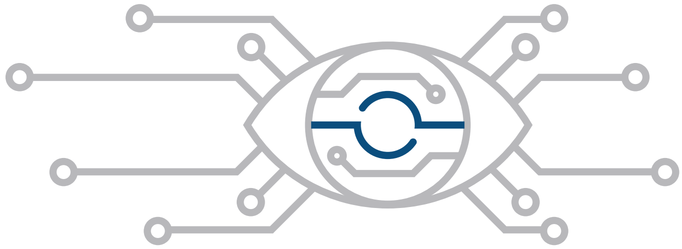

After decades of research and development, mostly confined to academia and projects in large organizations, artificial intelligence (AI) and machine learning (ML) are advancing into every corner of the modern enterprise, from chatbots to tractors, and financial markets to medical research. But companies are struggling to move from individual use cases to organization-wide adoption for several reasons, including inadequate or inappropriate data, talent gaps, unclear value propositions, and concerns about risk and responsibility.

This *MIT Technology Review Insights* report, commissioned by and produced in association with with JPMorgan Chase, draws from a survey of 300 executives and interviews with seven experts from finance, health care, academia, and technology to chart elements that are enablers and barriers on the journey to AI/ML deployment. The following are the report's key findings:

- **Businesses buy into AI/ML, but struggle to scale across the organization.** The vast majority (93%) of respondents have several experimental or in-use AI/ML projects, with larger companies likely to have greater deployment. A majority (82%) say ML investment will increase during the next 18 months, and closely tie AI and ML to revenue goals. Yet scaling is a major challenge, as is hiring skilled workers, finding appropriate use cases, and showing value.

- **Deployment success requires a talent and skills strategy.** The challenge goes further than attracting core data scientists. Firms need hybrid and translator talent to guide AI/ML design, testing, and governance, and a workforce strategy to ensure all users play a role in technology development. Competitive companies should offer clear opportunities, progression, and impacts for workers that set them apart. For the broader workforce, upskilling and engagement are key to support AI/ML innovations.

- **Centers of excellence (CoE) provide a foundation for broad deployment, balancing technology-sharing with tailored solutions.** Companies with mature capabilities, usually larger companies, tend to develop systems in-house. A CoE provides a hub-and-spoke model, with core ML consulting across divisions to develop widely deployable solutions alongside bespoke tools. ML teams should be incentivized to stay abreast of rapidly evolving AI/ML data science developments.

- **AI/ML governance requires robust model operations, including data transparency and provenance, regulatory foresight, and responsible AI.** The intersection of multiple automated systems can bring increased risk, such as cybersecurity issues, unlawful discrimination, and macro volatility, to advanced data science tools. Regulators and civil society groups are scrutinizing AI that affects citizens and governments, with special attention to systemically important sectors. Companies need a responsible AI strategy based on full data provenance, risk assessment, and checks and controls. This requires technical interventions, such as automated flagging for AI/ML model faults or risks, as well as social, cultural, and other business reforms.

Methodology

The survey that forms the basis of this report was conducted by MIT Technology Review Insights in September and October of 2022. Following are the key demographic details of the 300 executives who took part in it. The respondents hold senior technology roles in their organizations and are C-level executives: chief information officers; chief technology officers; chief data/analytics officers; and chief AI officers; as well as senior vice presidents or vice presidents of AI, of data platforms, or of engineering; and heads of AI and machine learning. These executives work in organizations that have between $1 billion and $500+ billion in revenue. The executives surveyed are located in four regions: Asia-Pacific, Europe and North Africa, Latin America, and North America. The industries represented in the survey include consumer goods and retail, financial services, hospitality and leisure, IT and telecommunications, manufacturing, media and marketing, pharma and health care, professional services, and transport and logistics.

"The journey is evolutionary, but the technology is revolutionary."

Samik Chandarana,
Chief Data and Analytics Officer for Corporate and Investment Banking, JPMorgan Chase

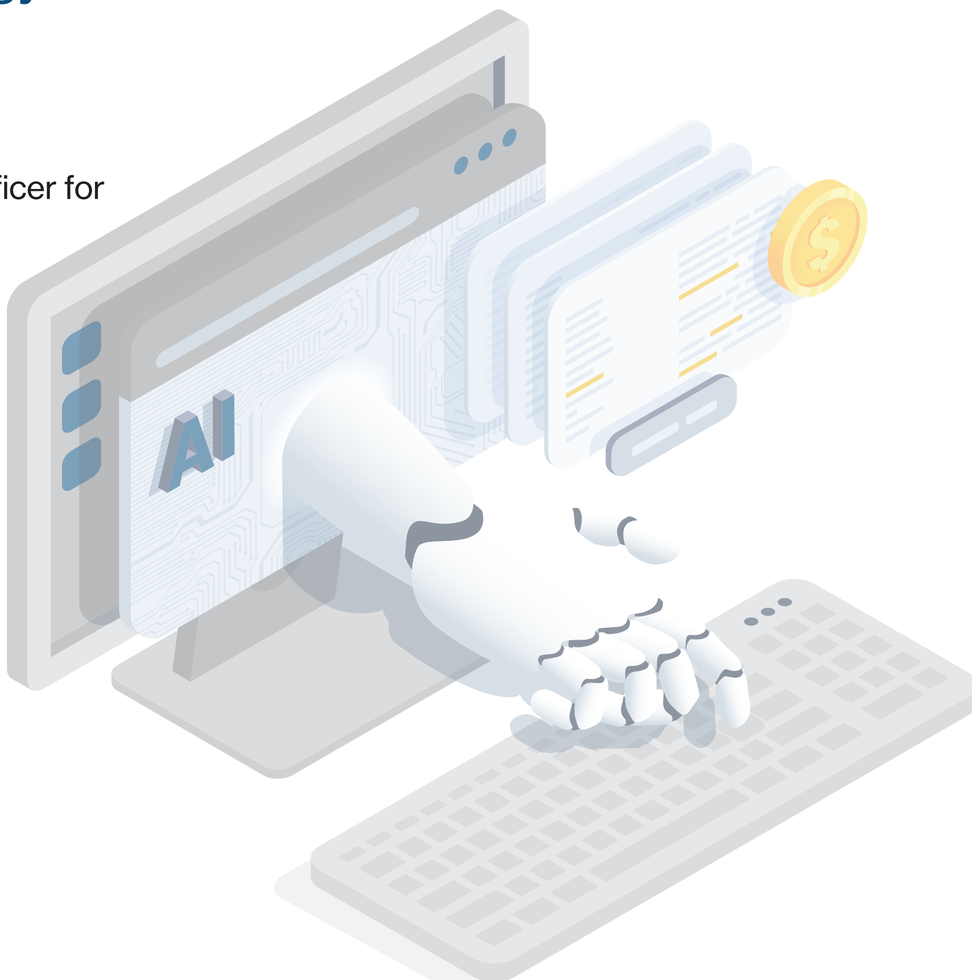

02 Machine learning is serious business

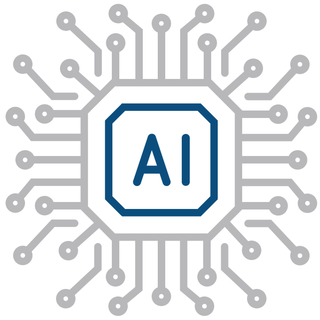

AI and ML are advancing in every corner of the modern enterprise. The achievements of AI/ML, like the protein-folding predictive powers of DeepMind's AlphaFold or the rapid, human-like responses of OpenAI's ChatGPT, garner media attention, but may also signal the rise of "boring" AI tools – supply chain optimization or automated form completion – that offer some of the clearest, if less exciting, indications of the technology's fitness for real-world use.

How AI will affect business depends on how these capabilities apply to vastly different domains. Prediction, for instance, is one of AI's greatest assets, according to Eric Boyd, corporate vice president of AI Platform at Microsoft. AI and ML can be useful for environmental and public health contexts like weather forecasting or infectious disease research, and in everyday contexts like anticipating customer behavior. Identifying anomalies in large datasets can detect fraud, cancer, and manufacturing defects.

A 2022 *MIT Technology Review Insights* survey, polling 300 respondents across four continents and nine industries, found that 71% of respondents use AI/ML across several projects, or heavily across the organization, with larger companies more likely to be active adopters.

Financial firms are not new to data science, but many powerful AI/ML techniques that were not computationally feasible before are making waves, says Samik Chandarana, chief data and analytics officer, corporate and investment banking at JPMorgan Chase.

Figure 1: AI/ML adoption

Seventy-one percent of all respondents surveyed say AI/ML is either heavily in use or in use across some projects.

27%	44%	22%	7%
MATURE/HEAVILY IN USE	IN USE FOR A FEW PROJECTS	EXPERIMENTAL	NASCENT

Source: MIT Technology Review Insights survey, 2022

Figure 2: AI/ML adoption among companies of different sizes

Larger companies, which tend to have more mature AL/ML capabilities and do more development in-house, are more likely to be active adopters of AI/ML

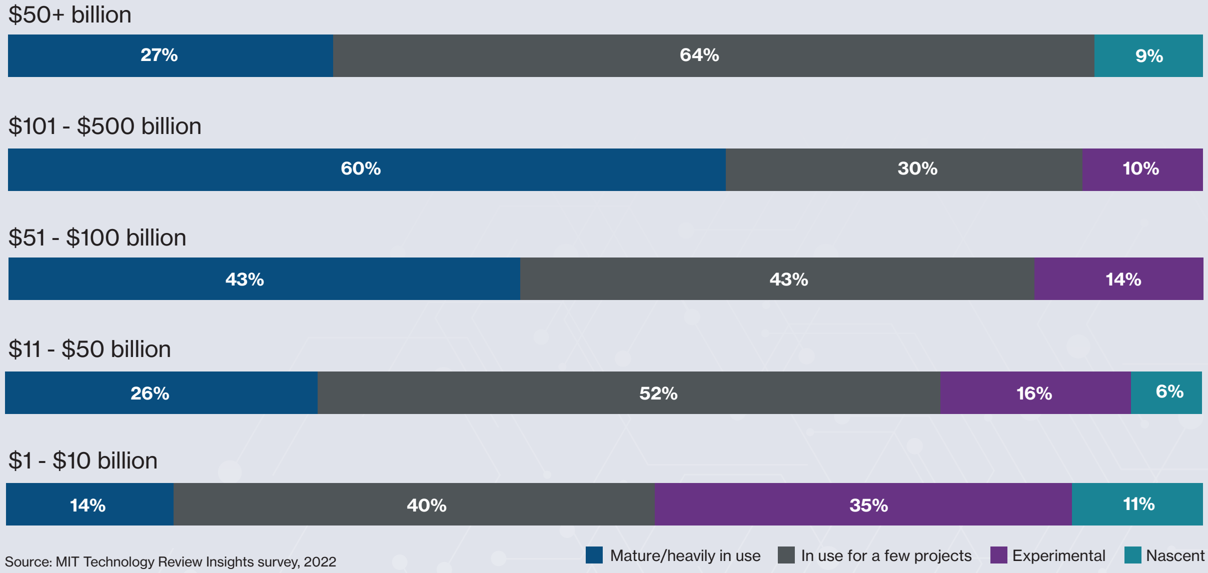

Source: MIT Technology Review Insights survey, 2022

“Every trader, banker, and payments person has always used data, so the journey is evolutionary, but the technology is revolutionary.” In finance, AI/ML can be used to optimize portfolios, underwrite loans, provide automated customer support, and detect fraud, as well as numerous other emerging use cases.[1,2]

As an example of some of the promising developments, JPMorgan Chase uses AI tools to flag companies likely to see activist shareholder pressure when filings drift from performance norms or expectations, and to navigate sanctions on Russian entities due to the war in Ukraine – names of sanctioned entities can be difficult to weed out from ordinary addresses without advanced techniques like AI. “We can reduce false positives,” explains Chandarana. “You can use ML to really power that ahead and increase people’s ability to get through that workload significantly faster.”

Human-like characteristics of AI

The future of AI/ML depends on how its capabilities can be adapted for different domains.

> “Monoliths—the idea that you can build a platform and it makes everyone happy—don’t work.”

Drew Cukor, Firmwide Head of AI/ML Transformation and Engagement, JPMorgan Chase

From proving to scaling

Companies often start AI/ML by identifying use cases and building models to surpass the status quo. Once proven, they move to other departments or functions, fine-tuning the tools, and securing engagement from the business.

JPMorgan Chase began exploring data science in the early 2010s, says Chandarana, but started working strategically with business units around 2017 through its Machine Learning Center of Excellence (MLCoE). Lidia Mangu, head of the MLCoE recalls, "There were a few pockets of data scientists and small teams doing ML on the side, but not as part of their formal role. There was no infrastructure for ML yet."

A turning point, Mangu says, came when the MLCoE convinced the credit business to provide data for predicting bond prices and used a deep learning-powered predictive model for the problem. "We showed them the results and they were very happy and said JPMorgan Chase already had the best bond pricing system in the market [at the time], so this baseline was really tough to beat!" Then the ML team identified other

Figure 3: Investment in AI use cases

Financial firms are not new to data science, but the capabilities at their disposal are novel thanks to new techniques and increased data availability.

AI'S HUMAN-LIKE CAPABILITIES	EXAMPLES
Generate new, engaging, realistic content and data based on existing ideas.	Personalized, interactive technology for customer service, generating and reviewing code, and drafting and sourcing complex documents.
Awareness and recognition of surroundings.	Cameras in self-driving cars.
Interpreting and ascribing meaning to information.	Converting hand-drawn sketches and notes to web pages.
Handling ambiguity.	Knowing whether someone searching online for an object wants to buy the object or be connected with fellow users.
Applying judgment.	Assessing if an unusual insurance claim is likely to be fraudulent.
Making decisions, even with incomplete information.	Weighing whether to follow an inconclusive health screening with a more invasive test, another screening, or an all-clear.

Source: Compiled by MIT Technology Review Insights based on data from Ana Landeta Echeberria, Artificial Intelligence for Business: Innovation, Tools and Practices, 2023[3]

AI/ML applications in health care

Advancements in AI technologies like federated learning and swarm learning are paving the way for novel AI applications in heavily regulated – and privacy forward – industries like health care.

One of the most promising early use cases of AI/ML in health care is in medical diagnostics. A 2022 study conducted by the U.S. Government Accountability Office revealed three emerging approaches for AI/ML diagnostic technologies: autonomous, adaptive, and consumer-oriented. From detecting cancer to predicting Alzheimer's risks, such advances help medical professionals identify and interpret the complex patterns of diagnostic data faster and more accurately, to improve patient treatment. The study also uncovered potential risks and limitations related to the efficacy of these advancements.

AI/ML technologies also are transforming medical research, such as helping to shorten drug development timelines from years into weeks or days.[4] Other advancements include:

- Alphabet's DeepMind and Meta recently deployed AI-powered protein folding prediction models[5], a task that eluded scientists for decades, with profoundly positive implications for drug development and life sciences.
- Generative AI tools like OpenAI's text-to-image DALL-E 2 are used by biotech researchers to design proteins never before seen in nature.[6]

	DESCRIPTION	POTENTIAL BENEFIT	POTENTIAL LIMITATIONS
Autonomous technologies	• Technologies that independently interpret images or other patient data to render a diagnosis.	• Fast, consistent information at the point of care. • Improved clinician capacity and patient access. • Earlier and more accurate detection.	• Developers may not be able to create and medical professionals may not adopt algorithms that diagnosis certain diseases autonomously.
Adaptive algorithms	• Technologies that update their algorithms by incorporating new patient data.	• May provide more accurate diagnoses or information by incorporating additional population or individual data. • Food and Drug Administration my be able to streamline its regulatory review of adaptive algorithms by reviewing potential changes to an algorithm during the initial review phase, rather than reviewing individual updates to algorithms. This could allow for rapid improvement of algorithms. • Could expand or improve features for users	• Changes in the algorithm data may lead to adverse outcomes such as inconsistent or poorer algorithmic performance.
Consumer-orientated technologies	• Technologies such as wearables and at-home devices that are marketed to consumers and may assist medical professionals in monitoring a patient's medical conditions.	• Can give medical professional more information about patients to improve diagnosis and treatment. • May increase access to care for consumers, particularly in underserved areas, such as rural settings, that lack specialists.	• Need further research to understand whether some devices improve patient outcomes. • Effectiveness may depend on patient's ability to understand or willingness to accept the health information presented.

Source: Compiled by MIT Technology Review Insights based on data from the U.S. Government Accountability Office, 2023.

potential use cases, using natural language processing models to extract information from support chats, news, and emails – a laborious process for humans, but easy for AI. "We did demo after demo and this led to an avalanche of new projects – up to 70 in the past year alone," says Mangu.

According to the survey, 28% of executives see scaling as their primary AI/ML deployment challenge. Other polls agree: McKinsey found only 15% of respondents successfully scaled automation across multiple parts of the business,[7] and Gartner data shows just over half (53%) of projects make it from AI prototype to production.[8]

Other key obstacles include hiring skilled workers, which survey respondents said was the number one challenge. Finding use cases, proving economic value, and outdated IT infrastructure were also primary issues flagged by respondents.

Aleksander Mądry, director of the MIT Center for Deployable Machine Learning, says data is a common scaling and deployment problem. "Sometimes companies cannot run a particular application because they do not have sufficient data, or it is difficult for the company to gather the data that is relevant," he explains. Data can be noisy or of poor quality, or fail to capture the trends needed for that solution.

73%

OF RESPONDENTS FROM COMPANIES $500+ BILLION IN SIZE SAY FINDING USE CASES IS THE BIGGEST OBSTACLE

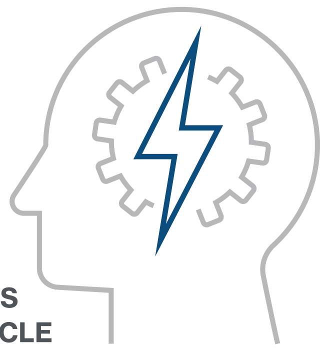

Models drift over time, he adds. "There are cases in the real world that you did not prepare your model for." He cites credit card fraud protection algorithms being disrupted during the pandemic, when consumer behaviors suddenly changed.

Mądry warns media hype has led to unrealistic expectations and uncritical enthusiasm from company leaders. "They don't actually stop and ask: is the problem I'm trying to solve an ML problem to begin with, do I have the right data for it, and is there even an objective?"

Figure 4: Key deployment challenges with AI/ML

Deployment challenges: two-thirds (66%) of respondents say hiring skilled workers is difficult. Question: What are the primary challenges your company has had with deploying machine learning? (Top three choices.)

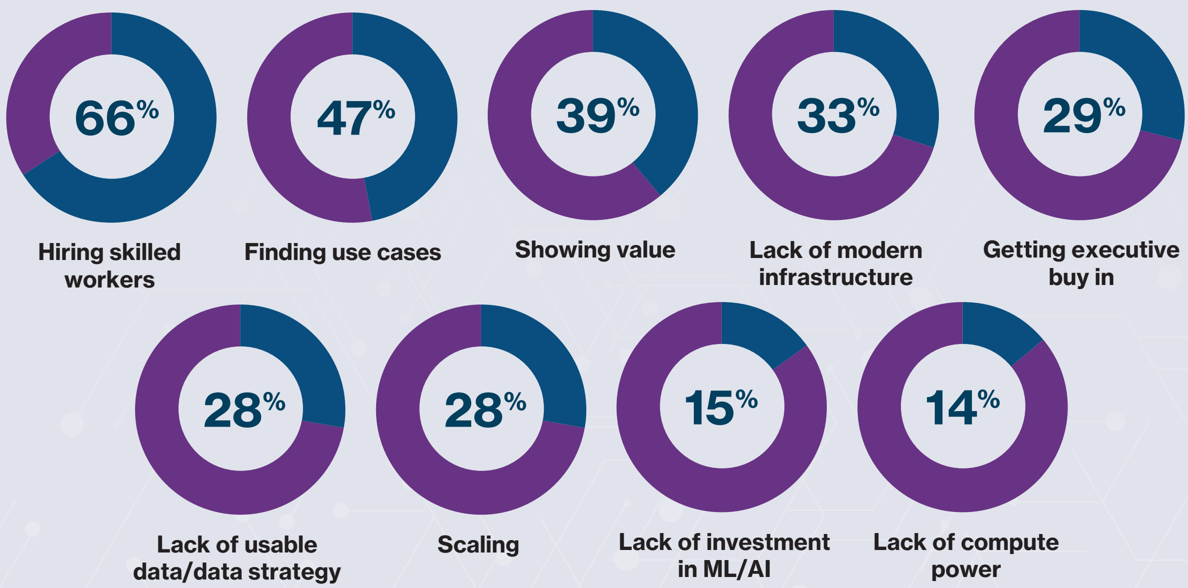

Source: MIT Technology Review Insights survey, 2022

03 Building out machine learning

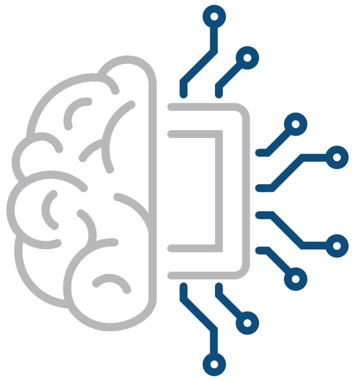

Although off-the-shelf AI/ML tools are available, companies that want to deploy at scale tend to develop internally. The survey found larger firms with mature AI/ML are more likely to build in-house.

CoEs are proving a popular structural approach to scaling AI and ML, and are used everywhere from Pfizer to Procter & Gamble.[9] "There will be commonalities across different use cases in the way that you want to brigade. You want this core expertise," says Mądry. "The center should have an internal consulting unit in which they are brought on to different use cases in the company and work closely with people to deploy and prioritize prototype solutions."

Boyd adds that CoEs can calibrate ML into a hub-and-spoke model. "We recommend companies start with a central team, but each business unit needs AI experts because they understand the products, the customer, the concerns, and the things they're going to need to do effectively."

> "They bring the deep business knowledge, we bring the heavy-duty ML knowledge."

Lidia Mangu, Managing Director, Head of Machine Learning Center of Excellence, JPMorgan Chase

Figure 5: Characteristics of mature AI use

Of respondents from companies with mature AI programs, most (73%) say AI usage is tied to revenue, most (70%) say AI usage is federated across the company, and nearly half (45%) say their AI tools are largely built in-house.

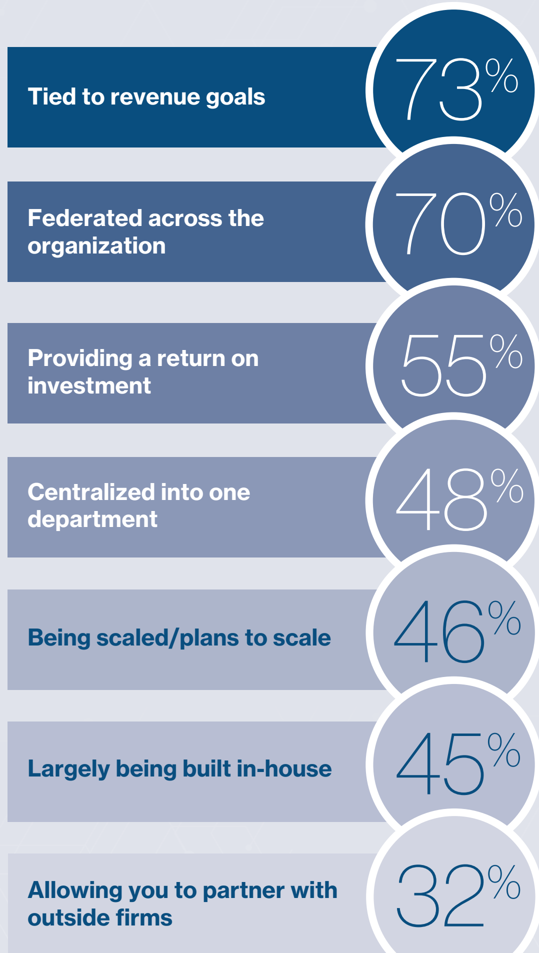

Source: MIT Technology Review Insights survey, 2022

> “Companies have to get their culture right so that they can embrace a data-driven and experimentation-driven mindset.”

Eric Boyd, Corporate Vice President of AI Platform, Microsoft

Drew Cukor, firmwide head of AI/ML transformation and engagement, at JPMorgan Chase, says it’s important to find the sweet spot between costly bespoke products that cannot be reused versus single platforms that cannot meet diverse needs. “Monoliths – the idea that you can build a platform and it makes everyone happy – don’t work,” says Cukor.

Cukor says the bank has three criteria for assessing whether an ML project is worthwhile. The first: “Reuse. In other words, when we [release it], can others grab it and reuse it so that similar use cases can quickly benefit from the investment.” The others are cost and speed – 120 days is an optimal timeline from project conception to production – a rough benchmark for a single project. This ensures these tools are useful across functions, from wholesale banking to retail. “There are certain models that we build for one asset class in markets for credit, but with customization, we can also use it for equities, rates, or other businesses,” says Mangu. The CoE capabilities Mangu’s team has built out to support the firm include natural language processing (NLP), speech analytics, recommendation systems, anomaly detection, time series analysis, reinforcement learning, and graph modeling.

There are parallels with the journey of Novartis, the Swiss pharmaceutical giant that is among the industry’s most active AI/ML enterprises. Its AI innovation lab researches and commercializes novel therapies, helping “answer the questions that matter to us and enrich the research around AI so that we can

Figure 6: Selected AI organizational models

STRUCTURE MODEL	FUNCTIONS
Star	• A center of excellence that works with all domains and departments. • Anyone can approach the team to get their AI model or application built. • Suitable for smaller companies, or companies just starting.
Matrix	• Dedicated to individual business problems. • Subset teams work directly with people in a domain. • Team members stay in their horizontal, share ideas, learn techniques, and compare domain-specific solutions. • Suitable for larger or AI-mature companies.
Fully embedded	• Every department has its own data scientists and engineers who build and deploy models, without reporting to a central team. • Enables focus and cross-functional communication, and eases incorporation of AI solutions. • Used by the largest and most AI-mature companies.

Source: Compiled by MIT Technology Review Insights based on data from Transforming Data with Intelligence (TWDI), 2021[10] enhance our drug discovery," says Iya Khalil, who until recently was the global head of the AI Innovation Lab at Novartis. "This is about scale; produce the algorithm once and now many different scientists can use it." The company has more than 200 data scientists working with biologists and clinical scientists.

In the finance sector, Mangu advocates partnership between business units and ML experts, but the latter must look outwardly and track field advances while also remaining focused on internal business needs. "They bring the deep business knowledge, we bring the heavy-duty ML knowledge, while we always try to keep up with the state of the art, and be aware of what's happening in the ML world," she explains. "If ML teams were only embedded in the business, they might have less time to keep current with ML advances."

Beyond the borders

Companies often lack the capacity to build AI/ML internally, and as a result regularly collaborate with vendors to meet their AI/ML needs. For example, Microsoft provides the engine for OpenAI's breakthrough ChatGPT and GPT family of models. And Microsoft's Azure platform offers an ML stack, explains Boyd – a base layer of developer tools for the ML lifecycle, a suite of cognitive services, and applied AI, which analyzes customers behavior. It offers not just computing power, but model-testing and scrutiny to ensure responsible development, says Boyd.

In the case of Novartis, a global pharmaceutical company, it worked with the tech industry to divide labor in ways that could play to its advantage. Novartis looks for fundamental AI research that tech companies have developed that it can extend in the health-care space. "That's one axis of the type of partnership that we want to have; an entity that is doing pure AI research, and we can use those models to drive our science," says Khalil.

Startups are also building technology and services to aid business adoption. Venture capital funds invested $67 billion in AI specialist firms in 2022.[11]

"When you deploy machine learning, it's not just about getting technical trust in the system as an engineer, it's also about the people who use them actually trusting it as well."

Aleksander Mądry, Director of MIT Center for Deployable Machine Learning and Cadence Design Systems Professor of Computing, MIT

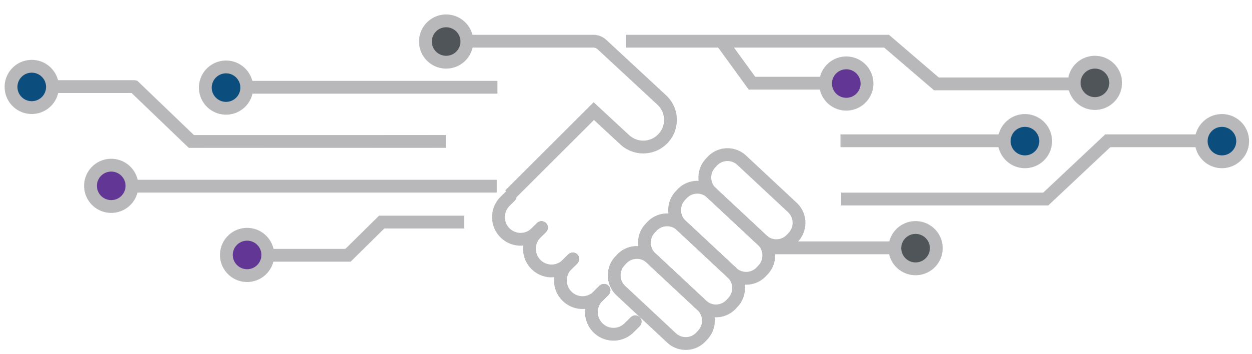

04 Mobilizing talent

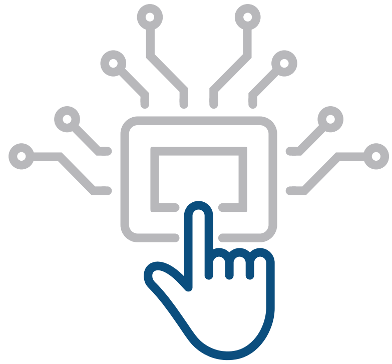

The battle for AI/ML talent is a test across the board; in a recent Deloitte study, one in four AI-experienced companies had a moderate skills gap, with a further 23% calling it major or significant.[12] The MIT Technology Review survey shows that hiring skilled workers is the single greatest challenge for respondents in all but the very largest firms. The talent problem has three dimensions: finding individuals with core AI/ML skills, cultivating blended skills that combine data science with domain knowledge, and producing a wider data-literate workforce.

Firms compete for a limited pool of ML talent, much of which is targeted by Big Tech firms. Mangu notes, firms are more likely to succeed in recruitment if they confidently articulate the unique opportunities they offer. "We have vast amounts of data, a very diverse range of use cases and projects, and the opportunity to make a large impact, to see internal and external clients using what you build," says Mangu. She adds that AI/ML teams at the bank are smaller than in a tech firm, meaning greater visibility and opportunities to progress.

Figure 7: Greatest challenges encountered by respondents while deploying AI/ML

REVENUE OF RESPONDENT COMPANY	GREATEST CHALLENGE	PERCENTAGE
$1-10 billion	HIRING SKILLED WORKERS	69%
$11-50 billion	HIRING SKILLED WORKERS	65%
$51-100 billion	HIRING SKILLED WORKERS	65%
$101-500 billion	HIRING SKILLED WORKERS	67%
$500+ billion	FINDING USE CASES	73%

Source: MIT Technology Review Insights survey, 2022

Figure 8: Cultural changes introduced by AI/ML adoption

Most respondents (73%) expect ML adoption will promote agility, learning, and innovation.

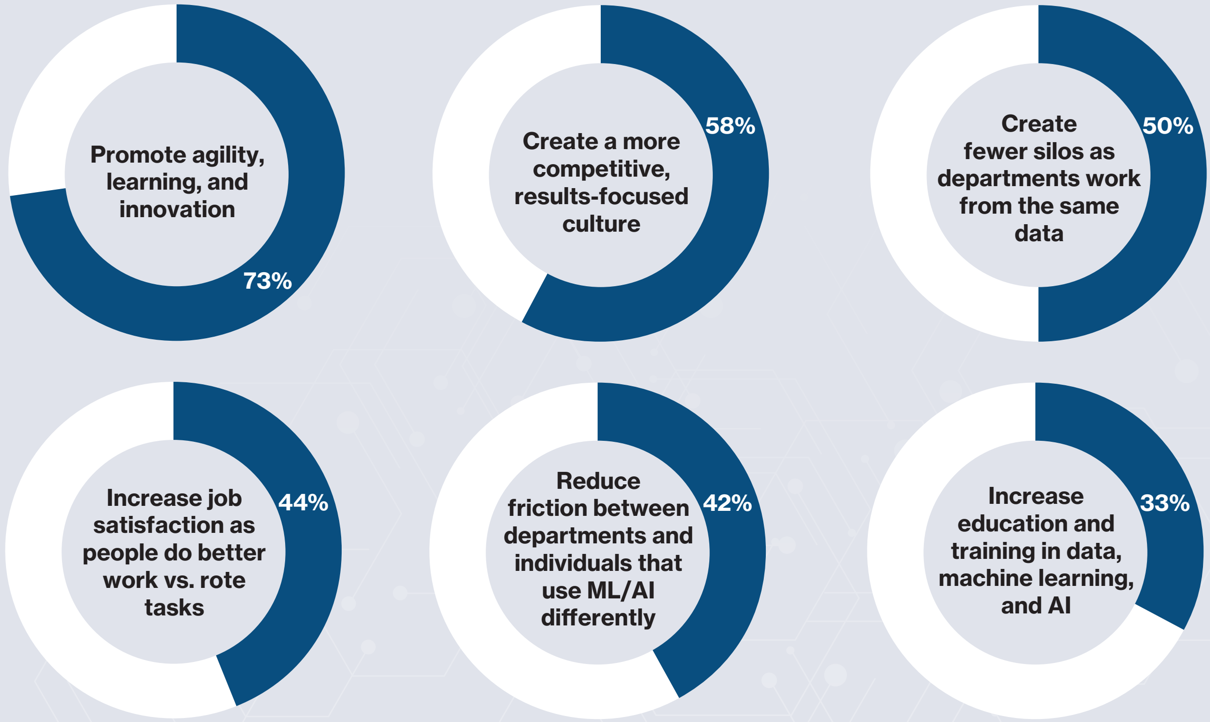

Source: MIT Technology Review Insights survey, 2022

Sought-after hybrid talent

Companies need AI/ML skills blended with domain knowledge. Novartis identifies three types of talent: AI specialists, hybrid talents, and scientists. "You want all three of those in your organization, and then you have to actively be able to connect the dots between these to drive progress," says Khalil.

Domain knowledge is vital to ensure AI/ML models appropriately account for the input data, respond to a clearly definable problem, and anticipate regulatory or governance risks. The AI for Science research community identifies domain knowledge as a "grand challenge" in developing AI.[13,14] This is seen in issues like applying AI to credit delinquency: ML models could improve predictive accuracy, reduce nonperforming loans, and even increase lending to creditworthy customers who lack credit history (which may reflect socioeconomic disadvantage rather than true risk). Because the number of credit delinquencies is so small and because encouraging expansion of lending requires relying on "missing" data, domain expertise is critical to determine the right modelling approaches and evaluations of fitness for purpose.[15]

Chandarana warns against technology advocates promising AI/ML as a universal solution. With so much input from outside financial services, context can be lost, he argues, which can lead to overcomplicated solutions and data misinterpretation. Domain knowledge makes sure experts can question or overrule a prediction or decision in vital areas like medicine, says Mądry.

An emerging talent cluster is AI "translators," people who can bridge the gap between business and technical teams. This role translates business needs into requirements and interprets AI results, guides change management and AI integration, and provides expertise in user experience design.[16]

Company case studies

Upskill to bridge the AI/ML skill gap

Our survey showed that hiring skilled talent is one of the biggest challenges for companies implementing AI technologies. Here are two companies taking an upskilling approach to help bridge that skill gap.

Levi Strauss & Co.
Levi Strauss & Co. introduced a companywide digital upskilling initiative with AI literacy courses for merchandise planners, coding education, and coursework for employee skills in ML, design thinking, and product management. In 2021, it introduced Machine Learning Bootcamp[17], immersive training in coding and ML. Senior management believes with staff engagement in AI/ML, they are ready for digital transformation.

McKinsey
McKinsey upskilled[18] more than 500 employees in AI/ML techniques and applications. As a company-wide initiative, McKinsey realized one-size-fits-all training wouldn't be effective, as talent skillsets varied widely across the organization. To make training accessible and successful, they developed a three-tier training program:

- AI aware: a beginner-level AI training program appropriate for anyone in the company to learn about AI applications for business functions and the importance of data for AI applications for everyday work.
- AI ready: a series of custom AI bootcamps for those with enough domain knowledge to spot AI opportunities. Internal experts run the bootcamps to demonstrate newer, broader AI technologies, and then provide online tools for continuous learning.
- AI capable: an intensive six-month course to upskill technical talent for such roles as AI UX expert, AI researcher, and machine-learning architect or engineer.

While recruiting and hiring skilled talent likely will be part of any successful AI/ML implementation, upskilling can not only make AI technologies more accessible across an organization, but can make implementation and deployment more successful. McKinsey notes, "no one understands how to get things done – and make them stick – in your organization better than your own people."

Improving access to ML

At scale, AI/ML impacts workflows as well as staff roles and responsibilities. At best, it removes tedious, repetitive, and dangerous tasks, freeing people for advanced, value-added, and rewarding work. But some perceive digital transformation as a threat.

In a case study at the Wharton Business School, for example, IT team members were perplexed by the shift to cloud computing, since managing the physical data center was part of their role. Wharton used a pro-change mindset to reorient members who, liberated from physical infrastructure, could now learn about and engage with the broader business.[19]

> "Companies have to get their culture right so that they can embrace a data-driven and experimentation-driven mindset."
>
> **Eric Boyd**, Corporate Vice President of AI Platform at Microsoft

Boyd says culture is the biggest impediment to AI/ML progress. "Companies have to get their culture right so that they can embrace a data-driven and experimentation-driven mindset." Successful firms achieve symmetrical AI culture, with harmony and balance assimilating AI versus asymmetrical organizations that fail to synchronize effort.[20] A Deloitte study agrees: companies with leading AI practices have cultural characteristics like cross-organizational collaboration and workforce optimism.[21]

However, most companies are not taking concrete steps to support the workforce for AI/ML. "Organizations are not paying enough attention to training," says Mądry. "It's not just the training of ML engineers, but the personnel on the ground." Training must be continually refreshed because of workforce churn. It is a big investment, he says, "but this is what it takes to deploy machine learning in a reliable and valuable way."

Professionals who are not experts in AI, like customer service representatives, are no less important. These customer-facing employees must understand what the ML tells them, and what it does not, Mądry says. In the MIT Technology Review survey, personalized customer experience is the number one benefit ML offers.

Engaging the workforce is not one-directional; staff should be part of technology development. "When you deploy machine learning, it's not just about getting technical trust in the system as an engineer, it's also about the people who use them actually trusting it as well," says Mądry. "This involves guidance and empowerment so that they can confirm the predictions, or intervene when they think something is wrong." Research shows people at all levels of an organization need to trust algorithms, and be able to make decisions to override them, rather than needing to consult a higher-up authority.[22]

> "We want to make sure that every piece of data has lineage, and we can trace it back to the origin and transformation of that data."
>
> **David Castillo**, Firmwide Head of AI/ML Technology and General Manager of AI/ML Product Line, JPMorgan Chase

05 ESG for the AI era

Many AI/ML stories focus on risks and pitfalls, such as biased decisions reflecting flaws in training data, unpredictable behaviors, and even flash crashes caused by algorithm-based automated trading. Moving forward with environmental, social and governance (ESG) issues shines light on responsible AI, the catch-all term for the principles, policies, tools, and processes to ensure AI systems are developed and operated in the service of good.[23]

The calls for explainability and accountability of AI/ML are governance challenges, especially in domains like credit or criminal justice, where parties have a legal right to dispute certain decisions. Reforms like the U.S. Algorithmic Accountability Act of 2022[24] and the EU's proposed AI legal framework[25] show growing attention from governments to ESG risks – and a move to complement or replace industry self-regulation.[26,27]

A 2022 PwC poll suggests that responsible AI is a priority for firms, with 98% of respondents saying they had plans to advance responsible AI in 2022, focusing on compliance with regulations and protecting AI systems from cyber threats.[28]

Responsible AI cannot be bolted on; it requires an end-to-end governance model and resourcing. Robust data provenance and visibility is fundamental; without it, companies struggle with decisions. "We want to make sure that every piece of data has lineage, and we can trace it back to the origin and transformation of that data. We are ensuring that data is being handled responsibly for ML every step of the way," says JPMorgan Chase's David Castillo, firmwide head of AI/ML technology and general manager of the AI/ML product line.

OmniAI, JPMorgan Chase's in-house platform, helps data scientists and engineers (machine learning and data) find data, transform the data, create features, and train models in a responsible manner. OmniAI enables AI to deploy at scale, standardizing processes and providing security and controls for working with confidential information.[29] "By addressing controls at that level, we're creating a very interesting and uniquely streamlined opportunity for performing machine learning across the entire machine learning development lifecycle [MDLC]," says Castillo. "This means we are being responsible across the entire spectrum of the MDLC."

Once AI/ML is designed, tested, and built, it requires continuous attention. Contemporary data science tools continually learn and adjust in response to new data and trends. However, this can lead to unexpected errors and the phenomenon of model drift.
To anticipate risks, companies adopt algorithmic impact assessments, based on systematic and structured risk assessments from the ESG sphere, to evaluate pathways and outcomes linked to new algorithms.[30]

Pressure to improve accuracy during the training phase of model development can lead to overfitting to a particular data set, which leaves the model vulnerable to errors after the model training process. "We want to

Increasing government attention to ESG

ESG in 2023: A selection of key ESG regulations

Australia
Beginning in 2024, it plans to phase-in required company-level data of carbon footprint, greenhouse gas (GHG) emissions, and climate risk from large businesses as part of annual financial reporting.

Canada
From 2024, some banks, insurance companies, and financial institutions must disclose how they assess and manage climate-related risks and opportunities, and identify how this impacts business, strategy, and financial planning. They also must detail ESG metrics and targets, including for Scope 1, 2, and 3 GHG emissions.

European Union
The EU is enacting ESG business regulations across nearly all economic sectors, most of which take effect between now and 2026. These include:

- Enhanced corporate financial disclosures for sustainability spending, a taxonomy for sustainability activities, and supply chain due diligence rules.
- Rules for specific products to boost circular ecosystems and sustainability.
- New laws for product labels, with penalties for greenwashing.
- "Women on Boards" requires women hold at least 40% of nonexecutive director positions at large companies by June 2026, or 33% applied across executive and non-executive positions.

India
The top 1,000 listed companies must produce a Business Responsibility Report (BRR) as a part of annual corporate reporting, detailing ESG metrics. Metrics align with UN Sustainable Development Goals (SDGs), and cover ethics, transparency and accountability around sustainable goods and services, workers, stakeholders, human rights, environmental protection and restoration, political influence, inclusivity and equitable development, and consumer protection.

Singapore
The Green Finance Action Plan, based on its UN 2030 net zero targets, creates ERM rules for banks, asset managers and insurers, a Green Taxonomy (targeting the energy, transport, and real estate industries) and green loan framework, and disclosure requirements for ESG funds. It is developing a blockchain-based system for ESG data, using a common standard.

United States
The U.S. Securities and Exchange Commission (SEC) plans to enshrine rules that standardize corporate ESG reporting with its upcoming ESG and Climate Disclosure Rules. Federal and state governments are creating various systems of ESG regulations, such as the Inflation Reduction Act (IRA), which is the largest investment in climate change mitigation in U.S. history. The broad legislation governs investment products with ESG features, GHG emissions disclosures for the largest federal contractors, and NASDAQ board diversity reporting.

United Kingdom
The Streamlined Energy and Carbon Reporting (SECR) is one of several new rules that require entities like large companies and investors and traders to complete yearly sustainability and climate reporting, including documenting energy use, calculating carbon footprints, and accounting for GHG emissions. The Sustainability Disclosure Requirements (SDR) targets greenwashing with investment label regulations, new disclosure requirements, and rules to limit marketing using "green" language.

Source: Compiled by MIT Technology Review Insights with data from Brightest, 2023[31] maximize performance on holdout data, and that's our proxy metric for initial model performance," explains Mądry. Confronting this pressure, he says, can be like teaching a student who schemes to pass the exam based on memorizing the specific answers at the expense of learning the material at a deeper level.

A constant feedback and assessment loop should "evaluate whether the system is doing what we intended it to do, or if it happens to optimize internally as a result of the process." Mądry's team is working on automated toolkits that recognize when model performance is decreasing, to synthesize inputs that can help improve modeling and correct for the problem of drift.

Gartner defines model operations (ModelOps or AI model operationalization) as "focused primarily on the governance and lifecycle management of a wide range of operationalized AI and decision models, including machine learning, knowledge graphs, rules, optimization, linguistic, and agent-based models."[32]

ModelOps can ensure end-to-end governance and lifecycle management of analytics, AI/ML, and decision models. This helps validate models in production with effects spanning IT, risk, compliance, business requirements, and operational efficiency.[33,34]

For Castillo, ModelOps, based on a similar lifecycle management approach in software, is distinct. "The controls that we have on our technology and the controls that we have on our data aren't necessarily the same," he explains. When JPMorgan Chase reviewed the controls on the ML development lifecycle, 70% of controls traditionally applied to software development were unnecessarily applied to ML, and there were gaps for ML issues not pertinent to software.

Cybersecurity is important since AI/ML systems require large data sets, greater connectivity, and greater digitization, which expand the risk surface. Castillo's team includes more than a dozen security architects looking at threat vectors for each ML product area.

ML tools are emerging to help companies which lack oceans of data, or which are unable to use the data which they do hold. Transfer learning, reusing a pretrained model in a new ML model and leveraging small data sets to deploy large models in a reliable way will impact the unfolding of ML, according to Mądry. Generalized knowledge can be used when two models perform similar tasks, to reduce resources and labeled data required to train new models. This is a significant to the evolution of ML and is increasingly used in development.[35]

In health care, approaches are emerging to allow AI/ML innovations without running afoul of regulatory constraints on data privacy. "Everything around patient records has to be HIPAA-compliant [a U.S. federal law protecting patient health information], but there are things that could help us take it to the next level, like federated learning and swarm learning," says Khalil. Swarm learning is decentralized ML combining edge computing, blockchain-based peer-to-peer networking, and coordination which can maintain confidentiality without the need for a central coordinator.[36] Also, policymakers are proposing regulatory sandboxes, like the European Commission's 2021 Artificial Intelligence Act[37] and associated incentives to encourage the safe testing of novel AI systems.

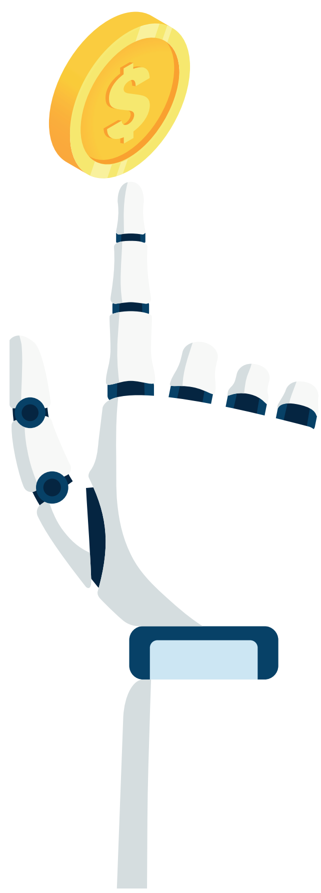

06 Conclusion

After waves of hope and disappointment, AI and ML are proving their utility in the business world, but scaling across the organization is a recurring challenge. Getting and wrangling the right data, identifying clear value propositions, and sourcing talent and skills are recurring challenges.

But these can be overcome by putting in place the right organizational structures, smart talent strategies that emphasize the unique opportunities of ML outside of the tech industry, and by harnessing the growing number of tech tools that can guide responsible AI and ModelOps. With the right strategy and sequencing, firms of all sizes can move forward AI/ML initiatives from projects and pilots to an organization-wide capability.

References

1. "Anti-Fraud Technology Benchmarking Report," Association of Certified Fraud Examiners and SAS, 2022, https://www.sas.com/content/dam/SAS/documents/briefs/third-party/en/acfe-anti-fraud-technology-112576.pdf.
2. Ana Landeta Echeberria, *Artificial Intelligence for Business: Innovation, Tools and Practices* (Springer Nature, 2022).
3. ibid.
4. Bill Briggs, "Novartis empowers scientists with AI to speed the discovery and development of breakthrough medicines," Microsoft, Nov 18, 2021, https://news.microsoft.com/transform/novartis-empowers-scientists-ai-speed-discovery-development-breakthrough-medicines/.
5. Ewen Callaway, "AlphaFold's new rival? Meta AI predicts shape of 600 million proteins," *Nature*, Nov 1, 2022, https://www.nature.com/articles/d41586-022-03539-1.
6. Will Douglas Heaven, "Biotech labs are using AI inspired by DALL-E to invent new drugs," *MIT Technology Review*, Dec 1, 2022, https://www.technologyreview.com/2022/12/01/1064023/biotech-labs-are-using-ai-inspired-by-dall-e-to-invent-new-drugs/.
7. Rohit Panikkar, Tamim Saleh, Maxime Szybowski, and Rob Whiteman, "Operationalizing machine learning in processes," McKinsey & Company, Sept 27, 2021, https://www.mckinsey.com/capabilities/operations/our-insights/operationalizing-machine-learning-in-processes.
8. "Gartner Identifies the Top Strategic Technology Trends for 2021" (press release), Gartner, Oct 19, 2020, https://www.gartner.com/en/newsroom/press-releases/2020-10-19-gartner-identifies-the-top-strategic-technology-trends-for-2021.
9. Thomas H. Davenport and Shivaji Dasgupta, "How to Set Up an AI Center of Excellence," *Harvard Business Review*, Jan 16, 2019, https://hbr.org/2019/01/how-to-set-up-an-ai-center-of-excellence.
10. Wilson Pang and Alyssa Rochwerger, "Choosing an Organizational Structure for Your AI Team," TWDI, May 3, 2021, https://tdwi.org/articles/2021/05/03/ppm-all-choosing-an-organizational-structure-for-your-ai-team.aspx.
11. "Artificial intelligence is permeating business at last: The age of "boring AI" will be anything but," *The Economist*, Dec 6, 2022, https://www.economist.com/business/2022/12/06/artificial-intelligence-is-permeating-business-at-last.
12. David Jarvis, "The AI talent shortage isn't over yet," Deloitte, Sept 30, 2020, https://www2.deloitte.com/us/en/insights/industry/technology/ai-talent-challenges-shortage.html.
13. "AI for Science: Mind the Gaps," NeurIPS AI for Science workshop, 2021, https://ai4sciencecommunity.github.io/neurips21.html.
14. Tirtharaj Dash, Sharad Chitlangia, Aditya Ahuja, and Ashwin Srinivasan, "A review of some techniques for inclusion of domain-knowledge into deep neural networks," *Nature*, Jan 20, 2022, https://www.nature.com/articles/s41598-021-04590-0.
15. "Domain Knowledge (Data Science)," CFI, Nov 2, 2022. https://corporatefinanceinstitute.com/resources/data-science/domain-knowledge-data-science/.
16. "Talent and workforce effects in the age of AI," Deloitte, 2020, https://www2.deloitte.com/content/dam/insights/us/articles/6546_talent-and-workforce-effects-in-the-age-of-ai/DI_Talent-and-workforce-effects-in-the-age-of-AI.pdf.
17. "Training Our Employees for a Digital Future," Levi Strauss & Co., May 17, 2021, https://www.levistrauss.com/2021/05/17/machine-learning-bootcamp/.
18. "Our own digital journey: Upskilling hundreds of McKinsey technologists in AI," McKinsey & Company, December 20, 2021, https://www.mckinsey.com/about-us/new-at-mckinsey-blog/ai-upskilling-for-over-500-firm-technologists.
19. "How Organizational Culture Shapes Digital Transformation," Knowledge at Wharton, Aug 9, 2019, https://knowledge.wharton.upenn.edu/article/fedex-digital-transformation/.
20. J. Mark Munoz and Oliver Degnan, "The Geometry of Organizational AI Culture," *California Management Review*, Jun 24, 2021, https://cmr.berkeley.edu/2021/06/the-geometry-of-organizational-ai-culture/.
21. "Fueling the AI transformation: Four key actions powering widespread value from AI, right now," Deloitte, October 2022, https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-state-of-ai-fifth-edition.pdf.
22. Brian McCarthy and Tamim Saleh, "Building the AI-Powered Organization," *Harvard Business Review*, July 2019, https://hbr.org/2019/07/building-the-ai-powered-organization.
23. Elizabeth M. Renieris, David Kiron, and Steven Mills, "To Be a Responsible AI Leader, Focus on Being Responsible," *MIT Sloan Management Review*, Sept 19, 2022, https://sloanreview.mit.edu/projects/to-be-a-responsible-ai-leader-focus-on-being-responsible/.
24. "H.R.6580: Algorithmic Accountability Act of 2022," Congress.gov, February 4, 2022, https://www.congress.gov/bill/117th-congress/house-bill/6580/text.
25. "Proposal for a Regulation laying down harmonised rules on artificial intelligence," European Commission, April 21, 2021, https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.
26. Luciano Floridi, "The End of an Era: from Self-Regulation to Hard Law for the Digital Industry," *Springer Philosophy & Technology*, Nov 10, 2021, https://link.springer.com/article/10.1007/s13347-021-00493-0.
27. Jakob Mökander, Prathm Juneja, David S. Watson, and Luciano Floridi, "The US Algorithmic Accountability Act of 2022 vs. The EU Artificial Intelligence Act: what can they learn from each other?" *Springer Minds and Machines*, Aug 18, 2022, https://link.springer.com/article/10.1007/s11023-022-09612-y.
28. "PwC 2022 AI Business Survey," PwC, 2022, https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-business-survey.html#protect-yourself-make-your-ai-responsible.
29. "Omni means "all"," JPMorgan Chase," 2023, https://www.jpmorgan.com/insights/technology/omni-ai.
30. "Examining the Black Box: Tools for assessing algorithmic systems," Ada Lovelace Institute and DataKind UK, 2020, https://www.adalovelaceinstitute.org/wp-content/uploads/2020/04/Ada-Lovelace-Institute-DataKind-UK-Examining-the-Black-Box-Report-2020.pdf.
31. "Key ESG Regulatory Policies, Laws, Rules, and Reporting Requirements in 2023," *Brightest*, January 27, 2023, https://www.brightest.io/esg-regulations.
32. "Information Technology Gartner Glossary: ModelOps," Gartner, 2023, https://www.gartner.com/en/information-technology/glossary/modelops.
33. "ModelOps vs. MLOps," ModelOp, 2022, https://www.modelop.com/modelops-and-mlops/.
34. "Top 10 Questions About MLOps and ModelOps," ModelOp, 2022, https://www.modelop.com/wp-content/uploads/2021/10/Top_10_Questions.pdf.
35. "Transfer Learning for Machine Learning," Seldon, Jun 29, 2021, https://www.seldon.io/transfer-learning.
36. Stefanie Warnat-Herresthal, et al., "Swarm learning for decentralized and confidential clinical machine learning," *Nature*, May 26, 2021, https://www.nature.com/articles/s41586-021-03583-3.
37. Tambiama Madiega and Anne Louise Van De Pol, "Artificial intelligence act and regulatory sandboxes," European Parliamentary Research Service, June 2022, https://www.europarl.europa.eu/RegData/etudes/BRIE/2022/733544/EPRS_BRI(2022)733544_EN.pdf

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review,* the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the U.S. and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. And through its growing MIT Technology Review Global Insights Panel, Insights has unparalleled access to senior-level executives, innovators, and entrepreneurs worldwide for surveys and in-depth interviews.

About JPMorgan Chase

JPMorgan Chase & Co. (NYSE: JPM) is a leading financial services firm based in the United States of America ("U.S."), with operations worldwide. JPMorgan Chase had $3.7 trillion in assets and $303 billion in stockholders' equity as of March 31, 2023. With over 55,000 technologists globally and an annual tech spend of $14 billion, JPMorgan Chase is dedicated to improving the design, analytics, development, coding, testing and application programming that goes into creating high quality software and new products. Under the J.P. Morgan and Chase brands, the Firm serves millions of customers in the U.S., and many of the world's most prominent corporate, institutional and government clients globally. Visit http://www.jpmorganchase.com/tech for more information.

JPMorgan Chase & Co.

Illustrations
Illustrations assembled by Tim Huxford with elements from Shutterstock.

MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com