MIT Technology Review Insights

Produced in association with

Adobe EY OWKIN

Our global poll examines key decision points for putting AI to use in the enterprise.

Generative AI deployment: Strategies for smooth scaling

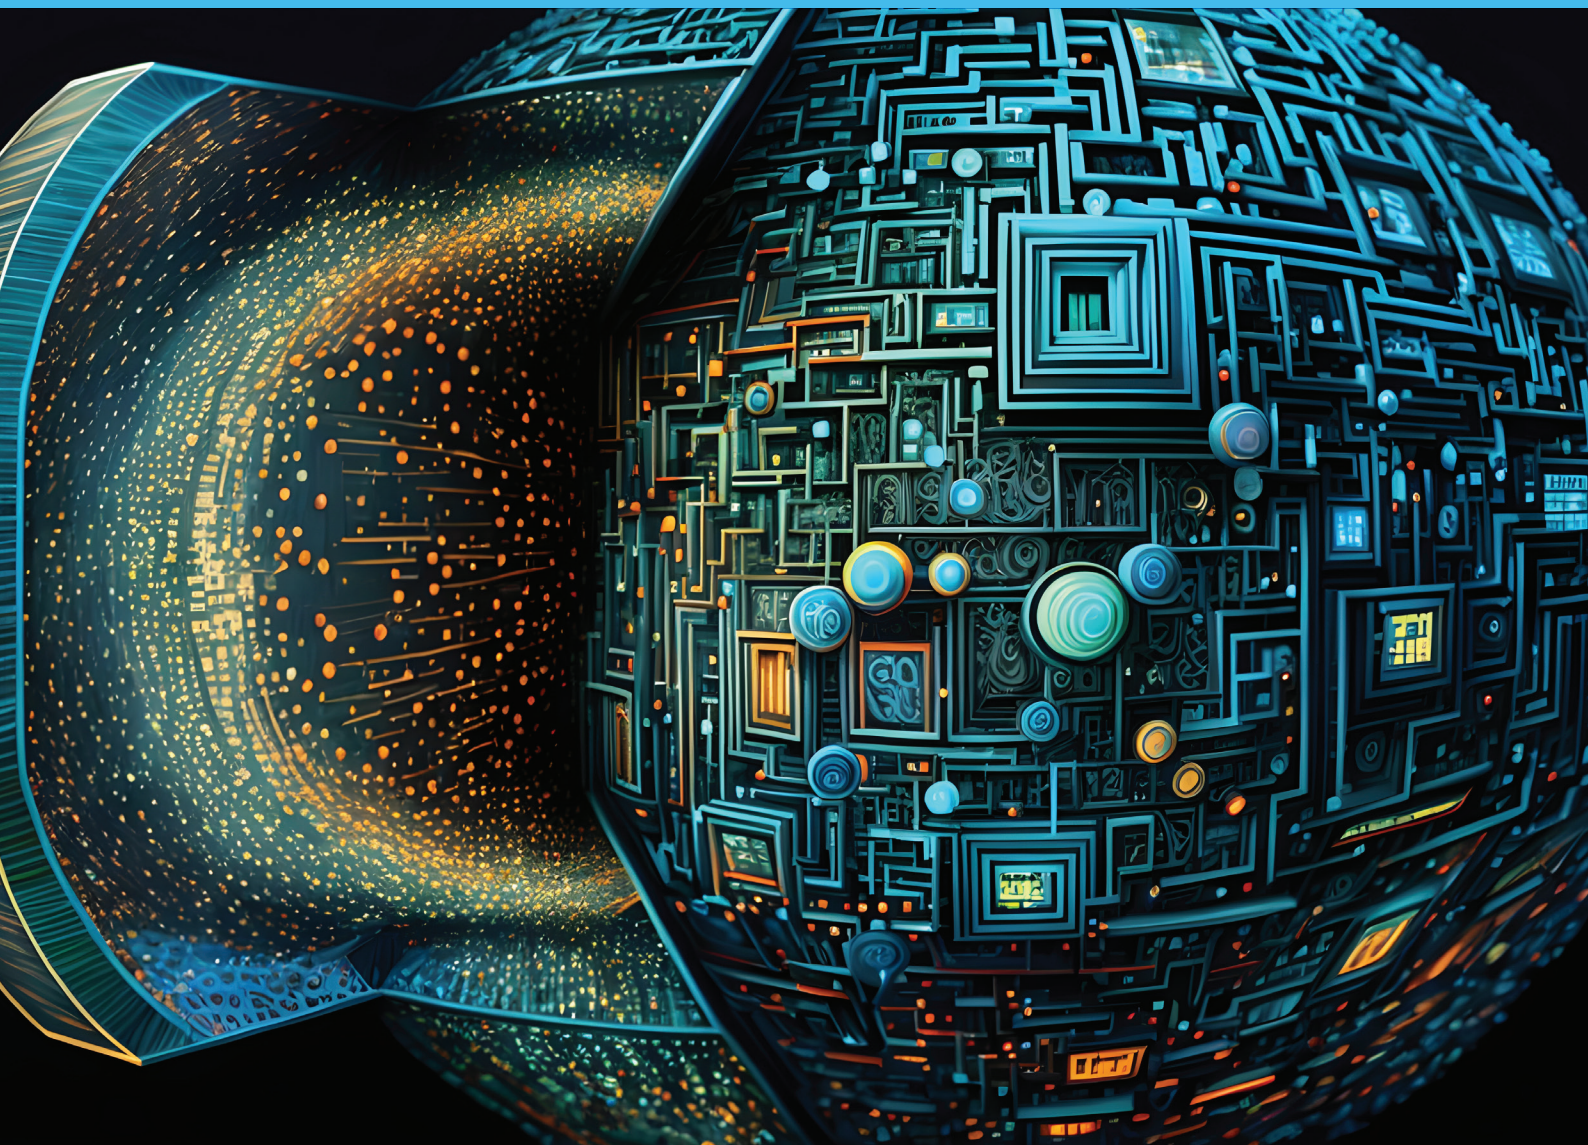

Preface

"Generative AI deployment: Strategies for smooth scaling" is an MIT Technology Review Insights report sponsored by Adobe, EY, and Owkin. To produce this report, MIT Technology Review Insights conducted a global poll of 1,000 executives. The report also draws on nine in-depth interviews with academics and data and AI experts, conducted in July and August 2023. Adam Green was the author of the report, Teresa Elsey 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 individuals for their time and insights:

Henry Ajder, AI expert and adviser

Alex Alonso, Chief Knowledge Officer, Society for Human Resource Management

Ben Armstrong, Executive Director and Research Scientist, Industrial Performance Center, MIT

Azeem Azhar, Founder, Exponential View

Steven Basart, Manager of Research and Development, Center for AI Safety

Ben Lorica, Data Scientist and host of The Data Exchange podcast

Thomas Kochan, Professor, MIT Sloan School of Management

Annemarie Schaeffer, Vice President of Research, Society for Human Resource Management

Julie Shah, Professor, Computer Science and Artificial Intelligence Lab, MIT

Poll methodology

In July and August 2023, MIT Technology Review Insights polled business leaders about their organizations' approaches to implementing generative AI technologies. The 1,000 respondents are C-suite executives (35%), VPs or directors (35%), or managers (30%). Respondents are distributed among 11 industries, including consumer goods and retail, financial services, manufacturing, and pharma and health care. The responses are global in scope, representing the Americas (22%), Europe (29%), Asia-Pacific (28%), and the Middle East and Africa (21%). Respondents' companies range in size, with global annual revenue (in U.S. dollars) ranging from less than $500 million (31%) to more than $10 billion (15%).

CONTENTS

01

Executive summary

After a procession of overhyped technologies like Web3, the metaverse, and blockchain, executives are bracing for the tidal wave of generative AI, a shift some consider to be on par with the advent of the internet or the desktop computer. But with power comes responsibility, and generative AI offers as much risk as reward. The technology is testing legal regimes in copyright and intellectual property, creating new cyber and data governance threats, and setting off automation anxiety in the workforce. Organizations need to move quickly to keep up with stakeholder expectations, yet they must proceed carefully to ensure they do not fall foul of regulations or ethical standards in areas like data privacy and bias. Operationally, enterprises need to reconfigure their workforce and forge partnerships with tech companies to design safe, effective, and reliable generative AI.

To gauge the thinking of business decision-makers at this crossroads, MIT Technology Review Insights polled 1,000 executives about their current and expected generative AI use cases, implementation barriers, technology strategies, and workforce planning. Combined with insights from an expert interview panel, this poll offers a view into today's major strategic considerations for generative AI, helping executives reason through the major decisions they are being called upon to make.

Key findings from the poll and interviews include the following:

• **Executives recognize the transformational potential of generative AI, but they are moving cautiously to deploy.** Nearly all firms believe generative AI will affect their business, with a mere 4% saying it will not affect them. But at this point, only 9% have fully deployed a generative AI use case in their organization. This figure is as low as 2% in the government sector, while financial services (17%) and IT (28%) are the most likely to have deployed a use case. The biggest hurdle to deployment is understanding generative AI risks, selected as a top-three challenge by 59% of respondents.

• **Companies will not go it alone: Partnerships with both startups and Big Tech will be critical to smooth scaling.** Most executives (75%) plan to work with partners to bring generative AI to their organization at scale, and very few (10%) consider partnering to be a top implementation challenge, suggesting that a strong ecosystem of providers and services is available for collaboration and co-creation. While Big Tech, as developers of generative AI models and purveyors of AI-enabled software, has an ecosystem advantage, startups enjoy advantages in several specialized niches. Executives are somewhat more likely to plan to team up with small AI-focused companies (43%) than large tech firms (32%).

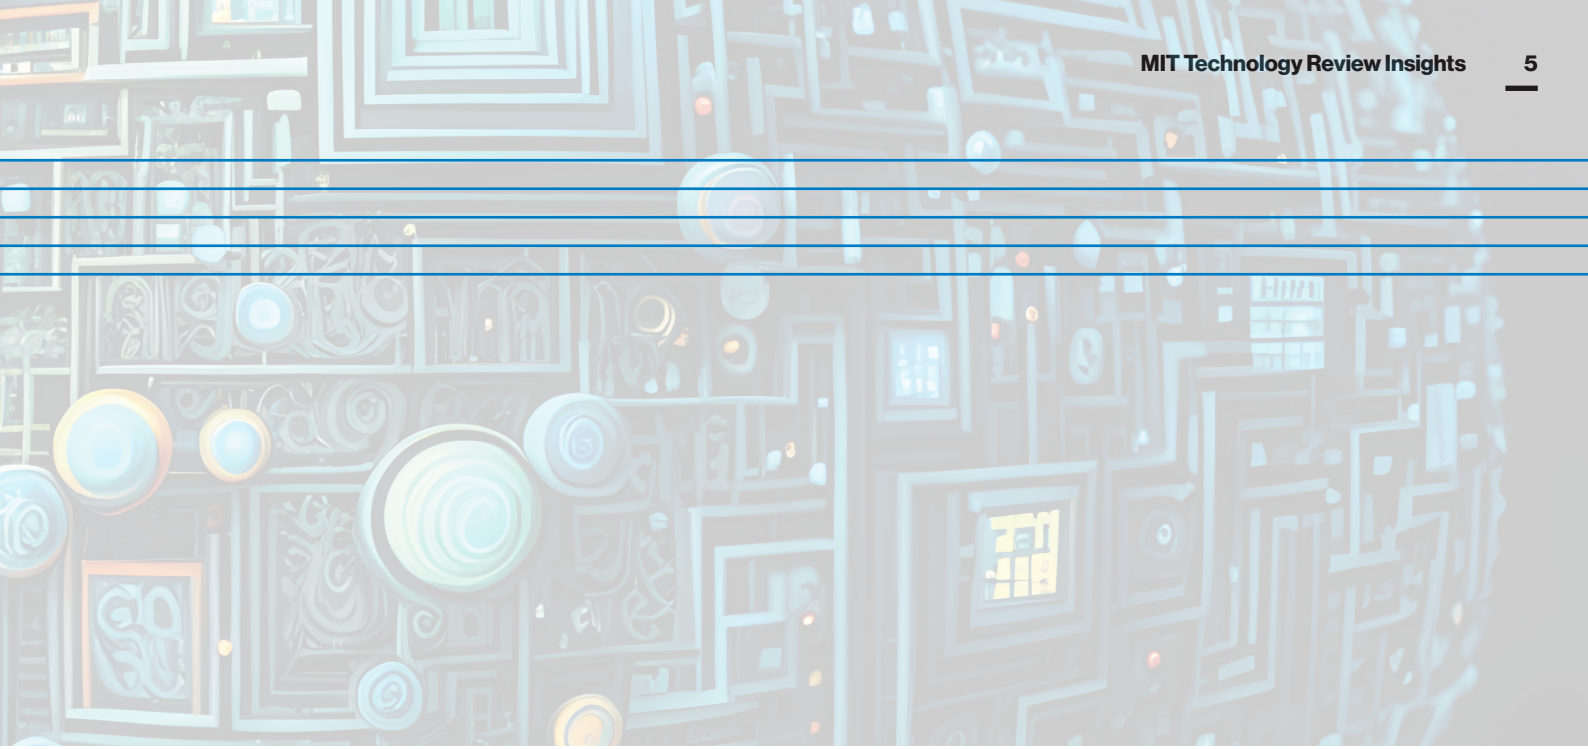

• **Access to generative AI will be democratized across the economy.** Company size has no bearing on a firm's likelihood to be experimenting with generative AI, our poll found. Small companies (those with annual revenue less than $500 million) were three times more likely than mid-sized firms ($500 million to $1 billion) to have already deployed a generative AI use case (13% versus 4%). In fact, these small companies had deployment and experimentation rates similar to those of the very largest companies (those with revenue greater than $10 billion). Affordable generative AI tools could boost smaller businesses in the same way as cloud computing, which granted companies access to tools and computational resources that would once have required huge financial investments in hardware and technical expertise.

• **One-quarter of respondents expect generative AI's primary effect to be a reduction in their workforce.** The figure was higher in industrial sectors like energy and utilities (43%), manufacturing (34%), and transport and logistics (31%). It was lowest in IT and telecommunications (7%). Overall, this is a modest figure compared to the more dystopian job replacement scenarios in circulation. Demand for skills is increasing in technical fields that focus on operationalizing AI models and in organizational and management positions tackling thorny topics including ethics and risk. AI is democratizing technical skills across the workforce in ways that could lead to new job opportunities and increased employee satisfaction. But experts caution that, if deployed poorly and without meaningful consultation, generative AI could degrade the qualitative experience of human work.

• **Regulation looms, but uncertainty is today's greatest challenge.** Generative AI has spurred a flurry of activity as legislators try to get their arms around the risks, but truly impactful regulation will move at the speed of government. In the meantime, many business leaders (40%) consider engaging with regulation or regulatory uncertainty a primary challenge of generative AI adoption. This varies greatly by industry, from a high of 54% in government to a low of 20% in IT and telecommunications.

02 Generative AI in the enterprise

Executives have endured several years of heavily hyped tech fads – including Web3, blockchain, and the metaverse – which have yet to find significant use cases and gain traction at scale. Generative AI, however, is different. By dint of its extraordinary domain flexibility and the many product layers now emerging to tune AI to specific needs, companies see few limits to its application. "Generative AI has more tangible utility than past tech hype cycles like crypto and metaverse," says AI expert and adviser Henry Ajder. "Though hype still exceeds current capabilities, adoption is real."

Generative AI broke into the public consciousness with the launch of OpenAI's ChatGPT in November 2022. But the large language models (LLMs) these tools are built on had been stealthily advancing for a decade, driven by improvements in hardware and technical breakthroughs such as 2017's Transformer model, a neural network that directs "attention" in ways that can effectively capture meaning and intent.[1] Since then, AI has captured the attention of enterprises, becoming an increasingly frequent mention on earnings calls (see Figure 1).

How fast are companies moving from pilots and prototypes to enterprise-wide deployment? What are the strategic considerations? An MIT Technology Review

Figure 1: Growing enterprise interest in AI
Share of U.S. companies mentioning AI on earnings calls.

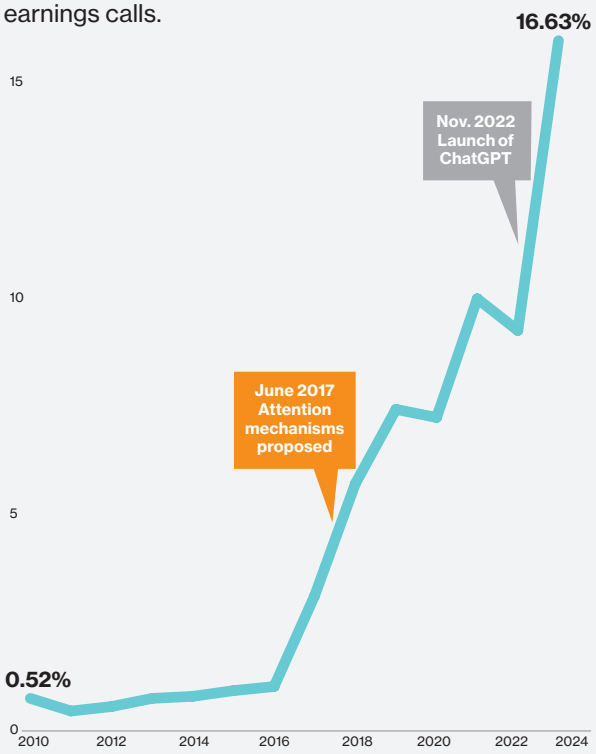

Source: Compiled by MIT Technology Review, based on data from "AI investment forecast to approach $200 billion globally by 2025," Goldman Sachs, 2023

Insights poll of 1,000 executives reveals that while nearly all companies consider generative AI a significant change in the technology landscape, with just 4% saying it will not affect their company, only 9% have fully deployed even a single use case across the organization (see Figure 2). This figure was as low as 2% in the government sector. IT and telecom (28%) and financial services (17%) businesses were the most likely to have deployed a generative AI use case.

Of course, there are significant technical, managerial, ethical, and legal hurdles to overcome on the way to implementing generative AI. Understanding the risks is the most common challenge, cited by 59% of firms (see Figure 3). This concern makes sense, as judges and courts are already presiding over a mounting file of legal challenges alleging the use of copyright material in the training of LLMs. "There's no way that these tools could exist or be brought into being without basically scraping the web," notes Ajder.

Legal risks for brands extend beyond IP and copyright to issues including liability and cyberattacks. As generative AI becomes "agentized" into, say, autonomous customer bots, companies could lose control of communications channels, predicts Ajder. "If you've baked a model into one of your customer-facing products, and suddenly it generates something defamatory or misleading and leads to actions which cause harm to individuals, or companies, you're going to be in real trouble." Cyber risks are mounting too. Hackers can use deepfake or synthetic media to deploy more authentic phishing

Figure 2: Generative AI deployment plans by industry
When will your organization begin deploying generative AI technologies (apart from isolated experiments)?

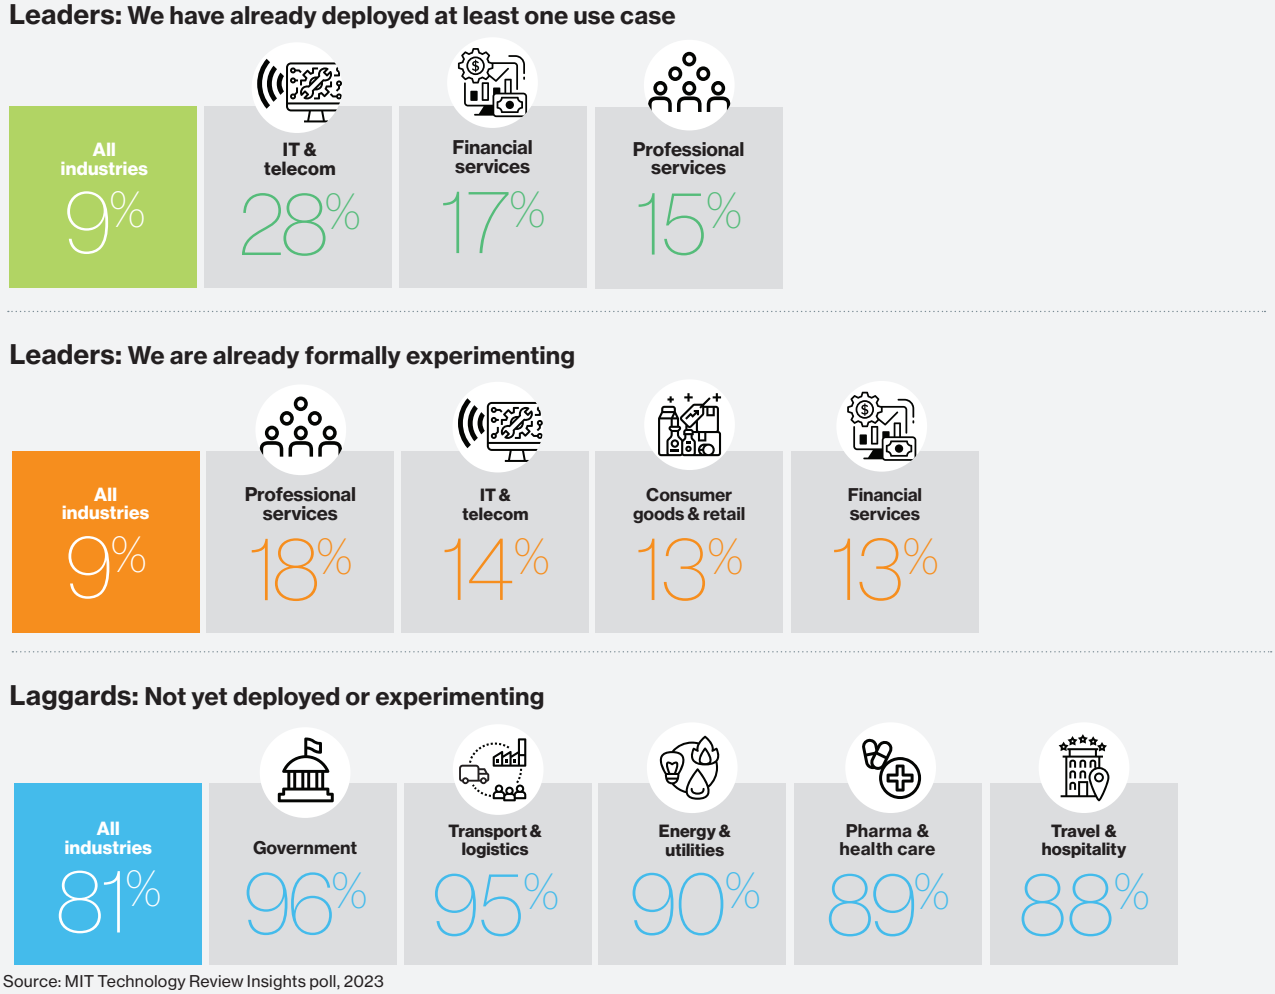

Source: MIT Technology Review Insights poll, 2023

campaigns, mimicking the communication style or voice of a known and trusted contact, whether a CFO or a supply chain partner.

The uncertain reliability of generative AI models is also a risk that is difficult to fully account for. The tendency of LLMs to hallucinate facts or deliver falsehoods is a significant obstacle in highly regulated or sensitive sectors like science or law. Generative AI skeptics like Gary Marcus, a cognitive neuroscientist, suggest that hallucinations may be an inescapable feature of generative AI models, since they are probabilistic engines rather than agents capable of true reasoning.[2]

The reliability problem could actually worsen as generative AI models exhaust the corpus of human-generated content on which they were trained, forcing AI to produce its own inputs leading, in the end, to degradation and model collapse.

The technical considerations executives must address to implement generative AI use cases are also not insignificant. Companies are simply not equipped to move from custom AI models to a "robust, production-grade system," says Ben Lorica, a data scientist and AI adviser. Many businesses are toying with simple models that search internal data and provide basic

Figure 3. Generative AI implementation challenges
What are the primary challenges your organization faces in successfully implementing generative AI? (Rank three.)

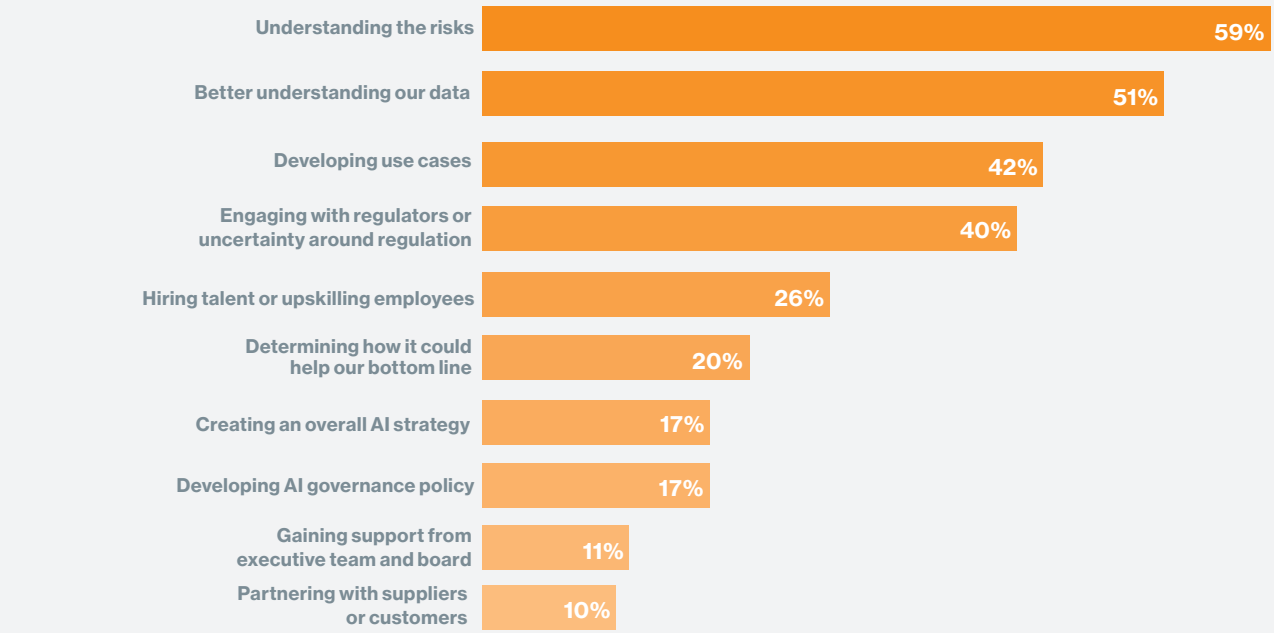

Source: MIT Technology Review Insights poll, 2023

"Generative AI has more tangible utility than past tech hype cycles like crypto and metaverse."

Henry Ajder, AI expert and adviser

conversational experiences, which is "an extremely narrow vision of what generative AI will ultimately be capable of," he says. "To fully realize the potential of this technology, there is still much work to be done in transforming prototypes into mature, enterprise-ready solutions."

While poll respondents are confident that generative AI will affect their businesses (only 4% say it will have no effect), the way it will affect them is less clear (see Figure 4). Executives project a wide range of positive and negative outcomes, from threatening their business offering (20%), to changing their product offering (18%), to allowing them to enter new sectors (17%).

The largest share (24%), however, say its primary effect will be to reduce their workforce. However, even as generative AI has been associated with doom-laden unemployment scenarios, our poll respondents and interviewees take a more measured view. They recognize the need to recalibrate the workforce and address automation anxiety.

Figure 4. Business effects of generative AI
How will generative AI affect your company? (Respondents selected one option.)

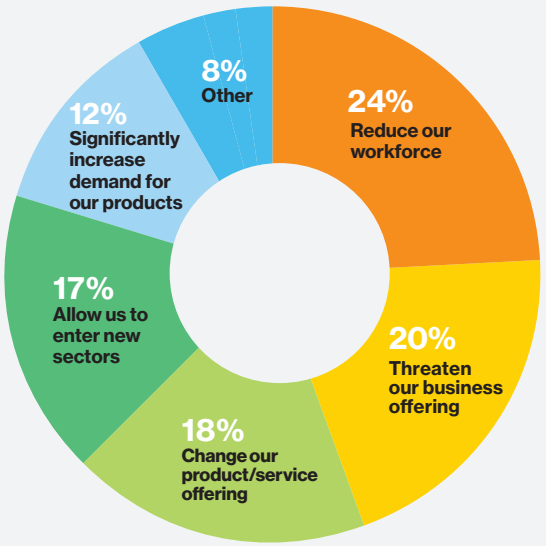

Source: MIT Technology Review Insights poll, 2023

Thomas Kochan, a professor at MIT's Sloan School of Management, says that, prior to generative AI, he was a skeptic about robots taking over human jobs. "I'm still not a believer that that's the ultimate outcome," he says, "but I believe that generative AI is a game changer. It is likely to have an impact across the workforce, from low-wage and lower-skilled employees, right up to the top of our occupation and skill ladders in the professions, in teaching, and in research."

Kochan worries not just about job displacement but also whether employees will have a say in, and share the fruits of, the automation wave. "Right now, workers don't have any explicit rights or opportunities to be involved in designing or helping to identify the problems on which these technologies can be used," he notes.

Overall, our poll respondents and interviewees are nearly unanimous in recognizing the potential for AI to revolutionize the enterprise. When we examine the speed with which they are deploying the technology within their own organizations, however, we see that their enthusiasm is tempered by substantial challenges and concerns, particularly around risk, and by uncertainty about the technology's ultimate effects.

03

Partnerships will prevail

When asked about their technology strategy, the vast majority of executives (75%) in our poll say they plan to partner with either a Big Tech company or a specialized startup to deploy generative AI (see Figure 5). Between the two, they are somewhat more likely to work with small providers (43% versus 32%). Very few companies (10%) consider partnering with suppliers or customers a top implementation challenge, suggesting that a rich AI service ecosystem has already emerged to support enterprises.

Given the number of technical and strategic considerations at play, it is unsurprising that only 19% of respondents say their organization is likely to go it alone. These respondents expect they will build their own generative AI technology and tools, either incorporating publicly available models (12%) or relying on proprietary or customized models (7%).

Big Tech companies are well-positioned to benefit from generative AI, with the likes of Google, Meta, and Microsoft making significant contributions to LLMs, and software giants like Salesforce able to quickly embed generative AI into their existing product ecosystems. But Big Tech does not have an unassailable advantage. Amazon Web Services (AWS) ditched its own model shortly after the launch of ChatGPT, while risk aversion meant that Google's experiments remained tentative and in-house until OpenAI forced its hand.[3] Also, tech transitions offer opportunities for new juggernauts to emerge, as Salesforce did during the enterprise transition to cloud, going from startup to major player.

Investors now expect generative AI to fuel a new generation of startups (see Figure 6). More than two-thirds of those surveyed by PitchBook (71%) believe generative AI will spawn a wave of technology unicorns (startups valued at more than $1 billion) over the next five years,[4] and an estimated $1.7 billion was generated across 46 deals in the first quarter of 2023 alone.[5] Niches for these startups include synthetic data

Figure 5. Primary technology strategies
When your organization begins deploying generative AI use cases, which is most likely to be its primary technology strategy? (Respondents selected one option.)

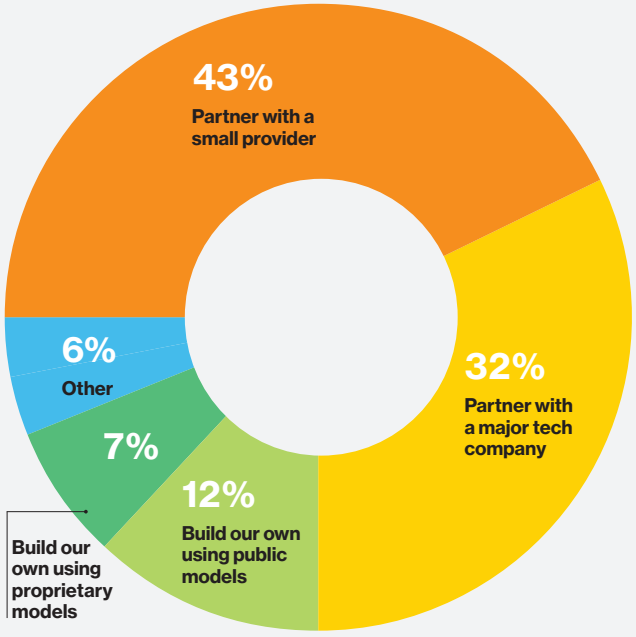

Source: MIT Technology Review Insights poll, 2023

"There's a layer of productization that's needed to deliver a useful product beyond the open API endpoints. Understanding that product process, understanding customer needs, creates an opportunity for startups."

Azeem Azhar, Founder, Exponential View

generation, training and feedback, model-tuning, and the development of chips optimized for training and running models.[6] Some firms are worried about relying on off-the-shelf platforms whose training data, models, and access terms could change at any time, which could drive them to work with smaller startups that offer products they could shape and direct.

"It isn't as straightforward as saying Microsoft and Google will do this," says technology commentator Azeem Azhar. "There's a layer of productization that's needed to deliver a useful product beyond the open API endpoints. Understanding that product process, understanding customer needs, creates an opportunity for startups." Azhar particularly sees opportunities for startups to develop tools for highly specific verticals and data sets, such as drug discovery.

Lorica sees opportunities for data-oriented startups and for those that can optimize AI infrastructure, improving these tools' energy efficiency and latency. Tools and services optimized for structured data and data warehouses will have to be remade to handle unstructured data. He also sees new business opportunities around deployment, or machine learning operations (MLOps), as organizations figure out how to wrangle gigantic generative AI models.

But the ground beneath startups is shifting, with heartache ahead for some. Each new update to a major LLM could, in theory, wipe out a problem that dozens or hundreds of companies were fixing – as was the case with OpenAI's release of GPT 3.5 Turbo, a tuning tool that lets developers customize models for specific use cases. Each LLM iteration can even quickly date its rivals. "There are things I might have to do in several queries in ChatGPT that I can do in a single query in Claude 2," says Azhar, referencing Anthropic's rival product. "So the question is whether improvements in core LLMs go and eat up that complexity." Ajder argues that many startups are just building wrappers around existing LLMs; niches like writing autocompletion, for example, are already thin and undifferentiated. "Startups are raising a lot of money," he says, "but many are just essentially dressing up other people's technology."

As it becomes clear that many fruitful generative AI implementations will involve partnership, business leaders will be wise to monitor this emerging AI landscape. When their organizations embark on generative AI use cases, they will have to not only figure out their own use cases and value proposition, but also find their place within a rapidly evolving ecosystem.

Figure 6. Investment in generative AI
Generative AI startup funding and deals have shot up in the first six months of 2023.

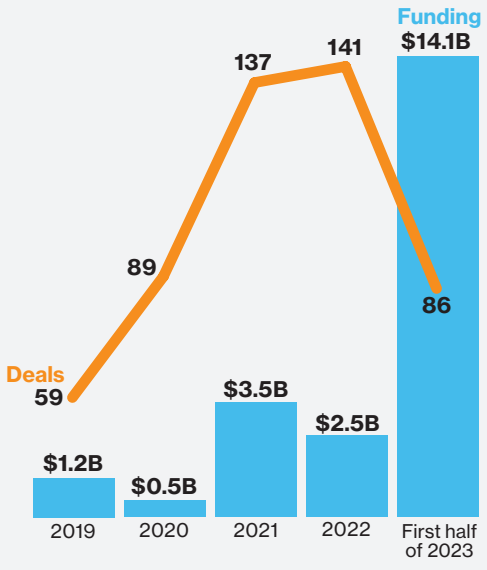

Source: Compiled by MIT Technology Review Insights, based on data from "The state of generative AI in 7 charts," CB Insights, 2023

AI for all: Democratizing the tech stack

The startup boom of the last decade was made possible, in part, by affordable enterprise technology, especially cloud computing and software-as-a-service (SaaS) solutions. These technologies gave companies of any size access to computation, applications, and analytics that once required large investments in technical personnel and servers. Generative AI is the next such wave of enterprise-enabling tech, allowing small firms to build new products and scale with ease. An AI feature that would once have taken a team of engineers years to build can now be developed in a few weeks, thanks to APIs provided by LLMs.

"Traditional forms of automation and AI required very significant resources to gain high performance," says MIT professor Julie Shah. "You need to structure the tasks, to cultivate and refine a high-quality data set, and to have high level, often PhD-level expertise. You would have a relatively inflexible technology considering the resources required, but you could achieve high performance. What is exciting about generative AI and LLMs is that they are flexible."

Smaller firms are seizing the opportunity. Our poll found company size has no bearing on the likelihood a firm is experimenting with generative AI (see Figure 7). Small companies (those with annual revenue less than $500 million) were actually more likely than mid-sized firms to have already deployed generative AI or to be formally experimenting with it – their poll responses were most similar to those of the largest firms (annual revenue more than $10 billion). These results suggest that mid-size firms could be caught in between, with more operational and brand risk at stake than smaller companies when they experiment, but with more limited capacity or infrastructure to advance safely toward firm-wide deployment than the largest companies.

Smaller firms were also the least likely to consider generative AI a threat to their business (14%, versus a poll average of 20%), and they were the least likely to cite risk (45% versus an average 59%) or uncertain regulation (31% versus an average 40%) as a primary challenge. Smaller companies were, however, more likely to struggle with creating and deploying an overarching AI strategy and less confident in ascertaining how generative AI could help their bottom line.

Figure 7. Generative AI deployment plans by company size (annual revenue)
When will your organization begin deploying generative AI technologies (apart from isolated experiments)?

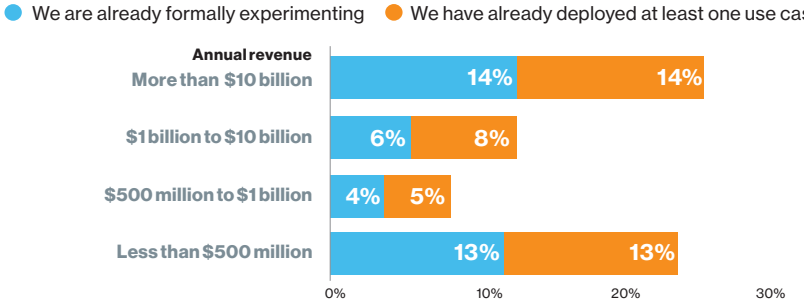

Source: MIT Technology Review Insights poll, 2023

"What is exciting about generative AI and LLMs is that they are flexible."

Julie Shah, Professor, Computer Science and Artificial Intelligence Lab, MIT

Owkin

Owkin is a TechBio company that combines the best of human and artificial intelligence to ensure every patient gets the right treatment. By understanding complex biology through AI, we identify new treatments, de-risk and accelerate clinical trials, and develop AI diagnostics.

Finding new drugs is getting harder every year.

Eroom's law (the inverse of Moore's law) is the observation that drug discovery is becoming slower and more expensive over time, despite significant technological advancement. The inflation-adjusted cost of bringing a new drug to market doubles roughly every nine years. Why? Two of the key reasons are a rise in regulatory standards and the incredibly complex biology of today's leading unmet disease areas (such as cancer and neuroscience).

But what if we could accurately simulate how a drug would work in a patient population before investing in a clinical trial? What if we could know whether a drug target – the biological bullseye and starting point for drug design – is worth pursuing, before paying hundreds of millions to validate it in the clinic?

Generative AI, adept at learning the patterns and structure of the data that it's trained on, gets its name from its ability to generate entirely new data, images, and text – the creation of new knowledge. This ability makes generative AI particularly useful for understanding human biology and helping scientists make sense of disease, illuminating the causal links between treatment, disease, and patient.

Foundation models (FM) are what generative AI is built on. FM are trained on massive amounts of unlabelled data – DNA sequencing for example – and power generative AI tools for a number of downstream tasks.

At Owkin, we have pioneered the development of FMs trained on patient data (such as images of cancer tissue samples from biopsies) to boost the performance of numerous AI tools for diagnosis and prognosis. This information enables health-care providers to offer tailored treatment plans, optimize patient outcomes, and reduce adverse reactions to medications. Building on this early success, we now use forms of generative AI to expedite the drug discovery process by finding new molecular targets against which to design drugs.

But drug discovery is only one piece of the puzzle. Drug development – successfully getting a drug through a clinical trial and approved by regulatory agencies – and diagnostics – precisely identifying a person's disease – are just as important. At Owkin, we believe that AI is instrumental throughout the drug development pipeline. We're partnered with Sanofi to identify drug targets, and with Bristol Myers Squibb to enhance clinical trial design and execution. We've also built an AI solution for digital pathology screening for colorectal cancer (MSIntuit CRC).

To reverse Eroom's Law to Moore's Law and reduce the difficulty of building new drugs, we need a new foundation to build upon. This foundation is being built right now, by companies like ours, and with the backing of the pharmaceutical industry – they, too, believe in the potential that this cutting-edge technology has to solve the unmet need of patients.

More generally, one can anticipate that, in the near future, novel FMs will be used by all scientists in biology to map experimental measurements at all scales of biology, from molecules to cells to tissues to patients, into actionable representations that will fuel discoveries and breakthrough moments at scale.

Jean Philippe Vert
Chief R&D Officer, Owkin

04 Working in the AI era

While generative AI is certain to have workforce effects, experts believe these will be more complicated, and possibly more positive, than some of the doom-laden estimates we have seen in recent months: The Organisation for Economic Co-operation and Development (OECD), for example, has estimated that nearly one-third of the workforce is at high risk, with highly skilled occupations the most exposed.[7,8] The consultancy Accenture projects that around 40% of working hours could be impacted by AI large language models.[9]

Our poll respondents' reactions seem more measured (see Figure 8). One-quarter of respondents expect the primary effect of generative AI on their business to be to reduce their workforce. That figure was much higher in industrial sectors like energy and utilities (43%), manufacturing (34%), and transport and logistics (31%). It was the lowest in IT and telecommunications (7%).

The experts we interviewed also offer more nuanced projections. Professor Julie Shah, who leads the Interactive Robotics Group of the Computer Science and Artificial Intelligence Laboratory at MIT, for instance, points out contextual subtleties that can limit automation's true human-replacing potential. "What is easy and hard for people is not the same as for machines," she says. "For instance, everybody can drive a car, but very few people can drive a commercial airplane. You would think that would make driving a car easier than flying an airplane for machines, whereas in fact, it's the exact opposite."

Figure 8. Projected workforce effects by industry

How will generative AI affect your company? (Respondents selected one option; percentage choosing "Reduce our workforce.")

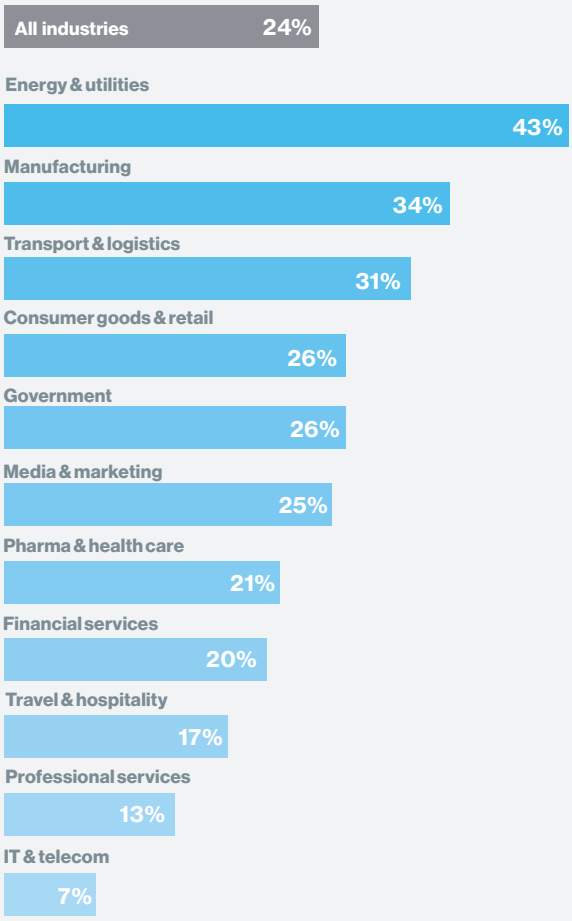

Source: MIT Technology Review Insights poll, 2023

Generative AI's reliability deficit will call for more, not less, human oversight. Ben Armstrong, executive director and research scientist at MIT's Industrial Performance Center explains, "Generative AI functions for workers as the inverse of a robot. Whereas robotic process automation excels in its robustness and consistency – it is really good at routine processes – it's really bad at flexibility." He continues, "Generative AI is the opposite, infinitely flexible in terms of text-based tasks, but not robust at all. It needs a human in the loop to be much more robust, which means this will be a tool but not a total automation solution."

According to the Society for Human Resource Management (SHRM), workers are not in a blind panic either. Fewer than 10% of Americans have been displaced or fear being displaced by AI, says Alex Alonso, the organization's chief knowledge officer. On the contrary, he says, 57% think they might enjoy their jobs more if their organizations provided them the opportunity to use and learn about generative AI.

Technical job postings are increasing in areas like next-generation software and cloud computing (see Figure 9), and in ethical AI. "For many companies, ethical just means legal. That may change," predicts Ajder. "We might start to see some meatier AI ethics roles, not people who are just coming up with ethereal principles, but going out and interrogating products, interrogating projects, not just for legal reasons, but for ethical reasons."

AI's contributions may be welcomed in sectors that are critically short of workers, such as manufacturing. "There are 700,000 job openings out there," observes Shah, of the U.S. economy. Far from replacing workers, generative AI could invigorate industrial sectors and help smaller firms compete in the labor market.

Technology-driven upskilling could emerge in surprising places. Alonso cites crane operators who, thanks to automation and drones, are stepping away from physically dangerous roles and shifting to analytical and creative areas, such as working with AI to design work schedules around weather patterns.

And if generative AI adoption follows a similar trajectory to robotics, it may both increase labor demand and make the companies that use it more attractive places to work. Armstrong cites evidence from several countries that

HR in the age of AI

Optimists think AI, used correctly, could improve fairness in recruitment, which can be rife with bias. It could, for instance, provide more sophisticated interview questions to gauge a candidate's true fit and reduce the information overload of hiring by providing interview summaries, approaches pioneered by recruitment tech startups like Metaview and Multiverse. These innovations could help ensure that hiring decisions are empirical and evidence-based, rather than relying on easy shortcuts like academic qualifications, which may not reflect true aptitude.[10] AI could also reduce the legal complexity of hiring international candidates, who might be overlooked because of the additional bureaucracy and paperwork.[11]

Annemarie Schaeffer, vice president of research at SHRM, says HR departments are already adopting generative AI in areas like content creation for job descriptions, recruiting, and developing salary benchmarks. "AI could support in any area of HR," she observes, "with the exclusion of that one-on-one empathetic kind of communications between HR and the employee base and executive leadership." Schaeffer thinks AI support for training and development could tackle burnout too. "Workers are burned out in large part because of the excessive and repetitive training that they need to conduct," she says. "That could easily be automated or helped along by generative AI."

MIT professor Thomas Kochan also sees an AI revolution in training. He argues for giving workers choices about the new skills and credentials they pursue, and then proactively considering them for new roles and opportunities. "It's about moving from an old-fashioned, top-down training model to a bottom-up learning model," he says, "where employees themselves make the decisions about how to get the skills they need to remain competitive in their own organization and in the external labor market."

when manufacturers adopt robots, on average, they also end up hiring more people. Shah's research has found that smaller companies use robotics not to replace workers but to attract fresh talent. "A young person would rather be working with the latest technology in a cleaner setting," Shah observes.

For the broader workforce, next-generation AI could alleviate burnout by automating repetitive, rote tasks, and boost engagement by enabling creative work. "The impact on knowledge work will be more qualitative than quantitative," says Lorica. "Work will be transformed rather than eliminated."

On the other hand, poorly executed automation can diminish humans' value and skills as they turn into machine appendages. Shah points out that while automation is pitched as an aid to humans, it can increase their cognitive load by requiring them to oversee, monitor, and fix its errors. In the worst case, it's harder for humans to find and fix an AI's errors than to do the work themselves. Shah also worries about anthropomorphized LLMs leading to misplaced trust in their instructions, especially in safety-critical jobs. Pilots, for instance, are more likely to accept an incorrect directive from an automated flight management system conveyed by a voice than via a text display.

Figure 9: Job postings growth between 2021 and 2022

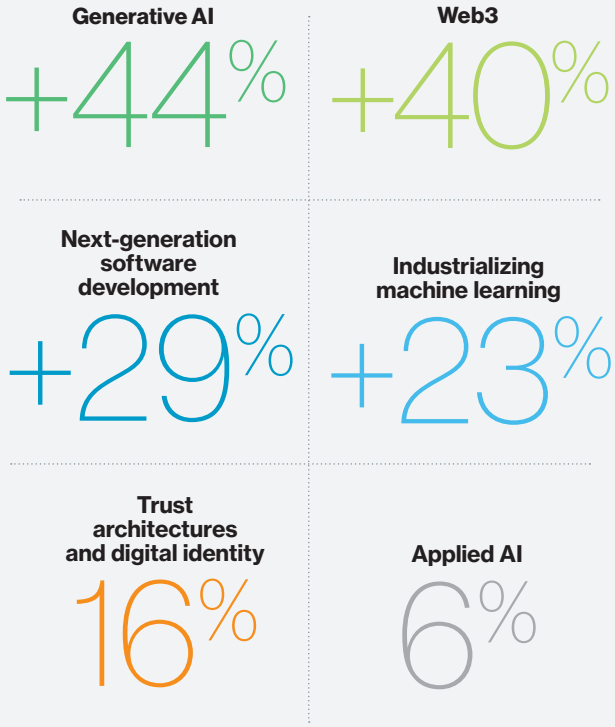

Source: Compiled by MIT Technology Review Insights, based on data from "McKinseyTechnology Trends Outlook 2023," McKinsey & Company, 2023

In any case, tomorrow's worker is more likely to become "AI-competent" than replaced, and AI could give entry-level or less traditionally educated employees the tools to perform more expert work. "The social norms will change, and it will be an expectation that you have competence in these tools like you do in Excel or Word," predicts Ajder. Alonso says the U.S. government has already taken a direct approach to reskilling its workforce and building "that prompt engineering mindset, that kind of general thinking around what makes a good AI worker."

> "Workers are burned out in large part because of the excessive and repetitive training that they need to conduct. That could easily be automated or helped along by generative AI."
>
> Annemarie Schaeffer, Vice President of Research, Society for Human Resource Management

Overall, our experts took pains to emphasize that technologies themselves do not determine workforce outcomes. "Technology is not exogenous; these are tools that through design, integration, and engineering can be used in different ways and have different consequences for work," says Armstrong. Choices that business leaders make now can lead to benefits for both organizations and their employees, particularly if they are based on a fuller appreciation of AI as a driver of innovation and value rather than a cost-cutting device and include employee input. "There's no law of physics that says this is the way a technology has to be applied," says Kochan. "Workers want a voice in those early-stage processes, not to block technological progress. We want a role in helping to shape technology use so that it both improves innovation and improves the performance of the organization."

Partner Perspective

Adobe

Over the last year, generative AI has become a central topic of conversation – inspiring questions about how and how much it will affect the way businesses operate, people work, and economies will change. This study found that most enterprises understand generative AI's potential to increase productivity and improve business processes, but they are early in adoption, as leaders assess safe and valuable ways to embrace this new technology. We'd like to share our experience.

A legacy of AI innovation

For more than a decade, Adobe has pioneered AI across our products, deploying hundreds of features to help customers to work smarter and faster. In March, we extended this legacy with Adobe Firefly, a new family of creative generative AI models designed to be commercially safe. We also announced generative AI capabilities to boost marketers' productivity and reimagine customer experience management.

For digital documents, Adobe Acrobat is the go-to PDF and e-signature solution for hundreds of millions of customers. With an estimated 3 trillion PDFs in circulation, generative AI can dramatically improve how everyone consumes, interacts with, and gets value from digital documents.

Innovating responsibly

Like all transformational technologies, generative AI comes with concerns and opportunities. And while automation changes the nature of work, this study confirms that most companies see generative AI as a means of helping people imagine and work faster and more creatively – regardless of where they work. And we agree.

Guided by Adobe's AI Ethics principles of transparency, authenticity, and accountability, we're building generative AI to enhance human creativity and productivity, deeply embedding it into customer workflows. We are pushing for industry-wide standards and helping establish best practices. We also advocate for Content Credentials, an open-source technology that serves as a "nutrition label" for images, providing transparency into whether images were created with AI.

When working with large language models (LLMs), we combine industry-leading solutions with strong safeguards to help ensure capabilities are deployed in a responsible way and our customers' data is safe.

Getting started with generative AI

While there is no single path to generative AI adoption, there are steps that can improve readiness and likelihood for success. Before deploying a single instance, establish protocols for evaluating and using generative AI. Identify the intersection of what the business needs and which tool – and partner – can best solve the problem or capitalize on the opportunity. Starting with focused projects and clear objectives and KPIs will help teams understand the value the technology is delivering, where to make changes, and where and when to apply more capabilities.

We believe generative AI can become a powerful force in helping everyone work more creatively and productively, make technology more accessible, and elevate human capacity for knowledge comprehension.

Whether an organization is making generative AI tools, using them or both, we all have a role in ensuring this exciting new technology benefits businesses and human beings. And while generative AI is still in its earliest stages, our customers are already using Adobe's new generative AI–powered tools to accomplish more, faster – and even change the definition of what's possible. We look forward to continuing to work with our customers, our employees, and industry partners to discover and build more value, together.

Randy Swineford
Senior Director of Product Management, Generative AI, Adobe

05

Governing in motion

Governments are busily drafting regulations, setting up task forces, and even creating new units of government to manage AI. This comes after years of legislative fury across the digital economy in data privacy, online safety, and competition policy. But the advent of generative AI seems to be spurring those years of debate into action, with 37 AI-related bills made law globally in 2022 (see Figure 10).

This flurry of activity will continue, but the most impactful legislation will take time to write and still longer to be put into effect. The European Union is taking a risk-based strategy, providing more stringent rules for more potentially hazardous uses of AI. But it is playing catch-up, with provisions in the EU AI Act, first proposed in 2021, rendered out of date by generative AI, such as a risk approach based on clearly defined use cases.[12] A group of 150 European companies, from Renault to Heineken, have criticized the text for harming European competitiveness, especially the attempt to regulate foundation models.[13]

> "Regulation and oversight is necessary to combat the inherent self-interest in the capital maximization that just naturally drives companies."
>
> Steven Basart, Manager of Research and Development, Center for AI Safety

Governments must consult widely to design robust and effective rules, and not just with the industry. "The car industry lobbied against seat belts, which now we almost take for granted, but they pushed very hard against something that would add a few cents to each vehicle, but goes against the profits," warns Steven Basart, manager of research and development at the Center for AI Safety. "Regulation and oversight is necessary to combat the inherent self-interest in the capital maximization that just naturally drives companies."

China has drafted some of the earliest and most detailed AI regulation, notably a 2021 ruling on recommendation algorithms, a 2022 regulation on synthetic or deepfake content, and a 2023 draft rule on generative AI.[14] Meanwhile, the U.S. is taking a light-touch approach, publishing AI blueprints, and rules for use in federal agencies, but no specific regulatory federal legislation yet.[15]

Smaller nations are vying for influence and seeking commercial opportunity through regulatory arbitrage. The UK, hosting a major summit in late 2023, is seeking a pro-innovation approach, similar to its open banking regulations that helped boost a generation of fintech startups.[16] Japan wants to kickstart an AI industry through a more liberal approach, such as replacing its copyright law to exclude IP used in training models.[17]

Comprehensive AI regulations are unlikely to be swift in passing. The UK Online Safety Bill has taken four years to pass from initial proposal to fully fleshed-out legislation. An instructive analogy might be the slow progress of the EU's General Data Protection Regulation (GDPR) – now considered a foundational personal data and information privacy law, and deeply globally impactful, it went into effect in 2018, more than six years after it was proposed and decades into the internet age.

Meanwhile, business leaders are deciding how to weigh these regulatory concerns, with some sharp differences between industries (see Figure 11). In our poll, 40% of respondents said they consider engaging with regulation or regulatory uncertainty a primary challenge of generative AI adoption, but this ranges from a high of 54% in government to a low of 20% in IT and telecommunications. Interestingly, some of the most regulated industries (such as energy and utilities and financial services) are less likely to consider regulation a challenge – suggesting perhaps that uncertainty about regulation is more concerning than the existence of regulation.

When it comes to governance, however, only 17% of respondents cite this as a top implementation concern. In our poll, industries that are less concerned about regulation tend to be more concerned with governance. IT and telecommunications is a prime example: respondents from that industry were both least likely to consider regulation a challenge (at 20%, only half the poll average) and most likely to name governance as a concern (28%). Perhaps there is an inverse relationship between industries' concern about externally imposed controls and their acknowledgment of their own responsibility to self-regulate.

Figure 10: Global AI-related bills passed into law

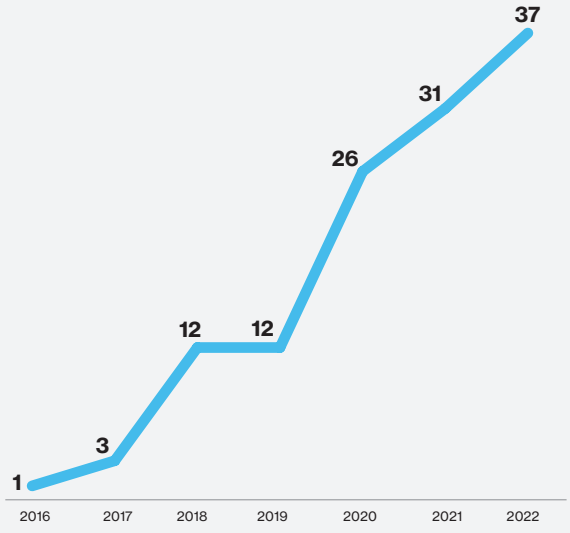

Source: Compiled by MIT Technology Review Insights, based on data from "The AI Index Report," Stanford Institute for Human-Centered Artificial Intelligence, 2023

Some of our interviewees are skeptical that AI can be regulated as a general-purpose technology, especially given the role of open-source software in its development and deployment, and they argue against comprehensive legislation. "How do you regulate open source? It's like trying to regulate mathematics," says Ajder. He believes that existing regulations could be sufficient to tackle some AI risks, and regulators should "only create new regulations where there is a genuine gap, rather than creating an entirely new framework that adds a whole layer of legal bureaucracy to navigate laws and harm landscapes as they stand."

Generative AI might prompt an overdue rethink of legal structures that were already fraying. Azhar, for instance, says copyright regulations are out of sync with the digital economy. "Copyright as a regulatory architecture has, for a very long time, not served creators or users, but the holders of IP, the publishers. Those who have followed the internet and digitization since the 1980s know there has always been this tension and the copyright system is deeply out of date," he says. "Will that get some form of overhaul? It would be a pity if we apply a very doctrinaire, harsh reading of copyright at this moment, because it was never designed for this kind of world."

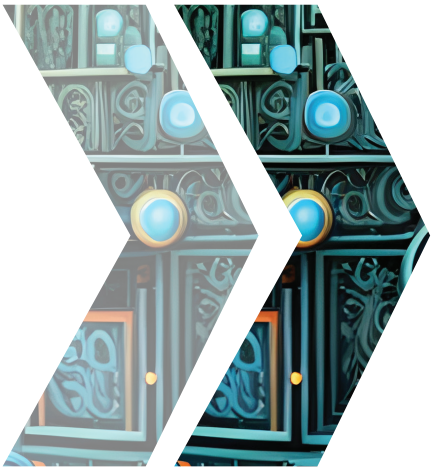

"It would be a pity if we apply a very doctrinaire, harsh reading of copyright at this moment, because it was never designed for this kind of world."

Azeem Azhar, Founder, Exponential View

Figure 11. Regulatory and governance challenges by industry
What are the primary challenges your organization faces in successfully implementing generative AI? (Respondents ranking challenge in their top three.)

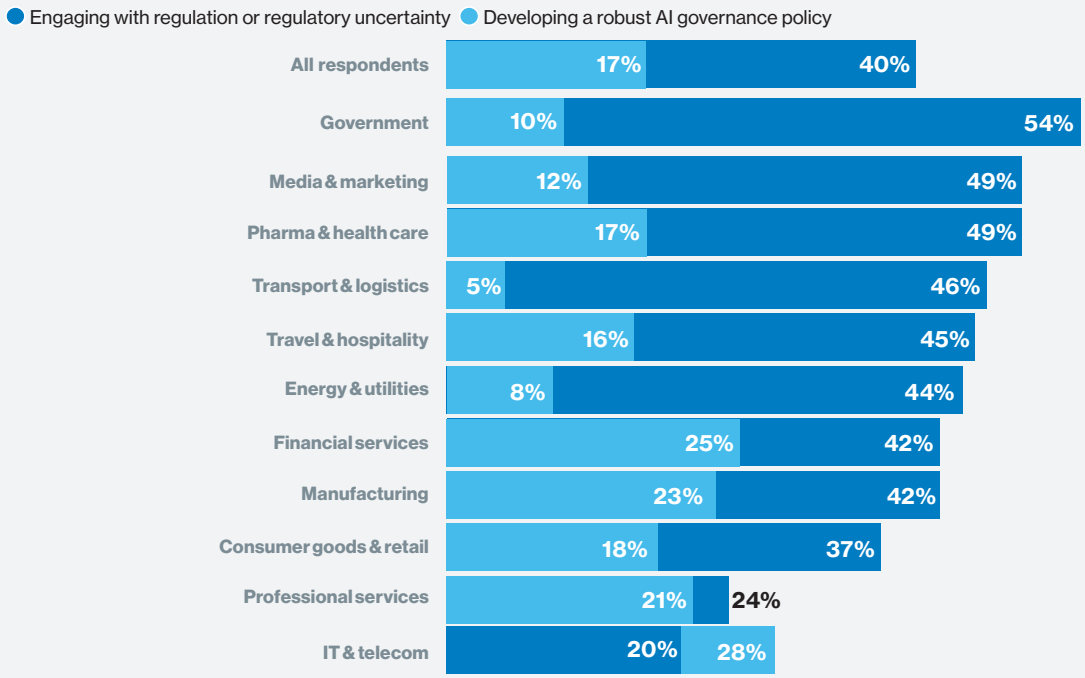

Source: MIT Technology Review Insights poll, 2023

EY

Artificial intelligence promises to transform economies and societies. At EY, our approach is designed to address a central challenge: If AI is a transformative force, how do we ensure it's a force for good?

We are addressing this challenge by investing $1.4 billion in AI and launching EY.ai, a unifying platform that brings together human capabilities and AI to help clients transform their businesses through confident and responsible AI adoption. Underpinning EY.ai – and everything we do in this space – are three guiding principles: (1) building confidence in AI; (2) creating exponential value through a holistic AI approach; and (3) augmenting human potential to drive extraordinary outcomes.

Building confidence in AI is an urgent need for our EY clients and the societies we all inhabit. Many are holding back from wider AI deployment because of a constraining uncertainty: AI is evolving faster than regulators can keep up. So a key component of the EY.ai platform is EY's Confidence Index, which EY teams use to help shape responsible AI guidelines and frameworks. We're working to bridge gaps between stakeholders and facilitate coordination across sectors, governments, and countries. As a global leader in high-quality audits, we've helped build confidence in financial markets, and we're committed to doing the same for AI.

Unlocking exponential value is critical for thriving in an AI-enabled age. This requires moving beyond piecemeal pilots to a broad rethink of strategy and operations. We're helping clients make the shift with the EY.ai Maturity Model, which benchmarks clients' state of AI adoption and empowers them to develop a strategic roadmap for harnessing AI's transformative potential. Meanwhile, EY's Value Accelerator helps clients prioritize the initiatives with the most potential to add value. Our professionals are helping clients identify areas across the business where AI can help enable value creation – from finance, legal, and tax, to risk, customer service, supply chain, and beyond.

Central to all of this is the human component. AI's best use is when it puts humans at the center and augments human potential. EY teams are helping clients automate mundane tasks with the goal of freeing employee time for higher cognitive work and meaningful interactions with customers and colleagues. This approach helps clients use AI to innovate job roles, empower their workforce, and scale collective intelligence.

As an organization built on human capital, we are actively working to augment human potential in EY people as well. Following an initial pilot with 4,200 EY technology-focused team members, the EY global organization is releasing a secure large language model called EY.ai EYQ. In parallel, we are upskilling EY people for the era of AI with a broad learning program to elevate AI skills.

Through this journey, we remain anchored to our purpose of building a better working world, guided by core values of integrity, respect, and ethics. By adopting a responsible approach to AI's design, deployment, and governance, we are helping harness AI's potential as a transformative force – and a force for good.

Nicola Morini Bianzino
Global Technology Officer, EY

06

Conclusion: Strategies for generative AI deployment

Generative AI is the most broadly consequential technology shift since the internet, and it is moving at a speed unlike any past tech capability. As firms now look to move from pilots and rudimentary applications to enterprise-wide deployment, they are developing strategies for tech partnerships and workforce redesign, while looking to governments for signs of the new regulatory architecture. This report points to strategic best practices for smooth scaling.

• **Pace for success.** Firms should move steadily and cautiously to enterprise-wide AI deployment, especially in light of evolving regulatory and legal regimes. "If you build a generative AI foundation and start deploying in a way where you scale your business on certain foundations, in six months' time, your entire business strategy could become legal liability," says Ajder. "You're going to have to dig up those foundations, at more cost than if you'd just thoughtfully implemented something narrower or had consulted further the likely direction of legislation." History shows that new technology adoption is often slowest at the start but can accelerate

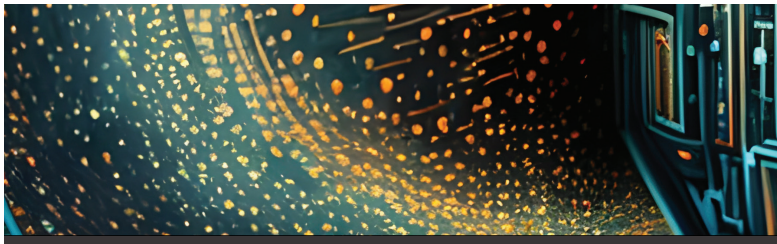

"If you build a generative AI foundation and start deploying in a way where you scale your business on certain foundations, in six months' time, your entire business strategy could become legal liability."

Henry Ajder, AI expert and adviser

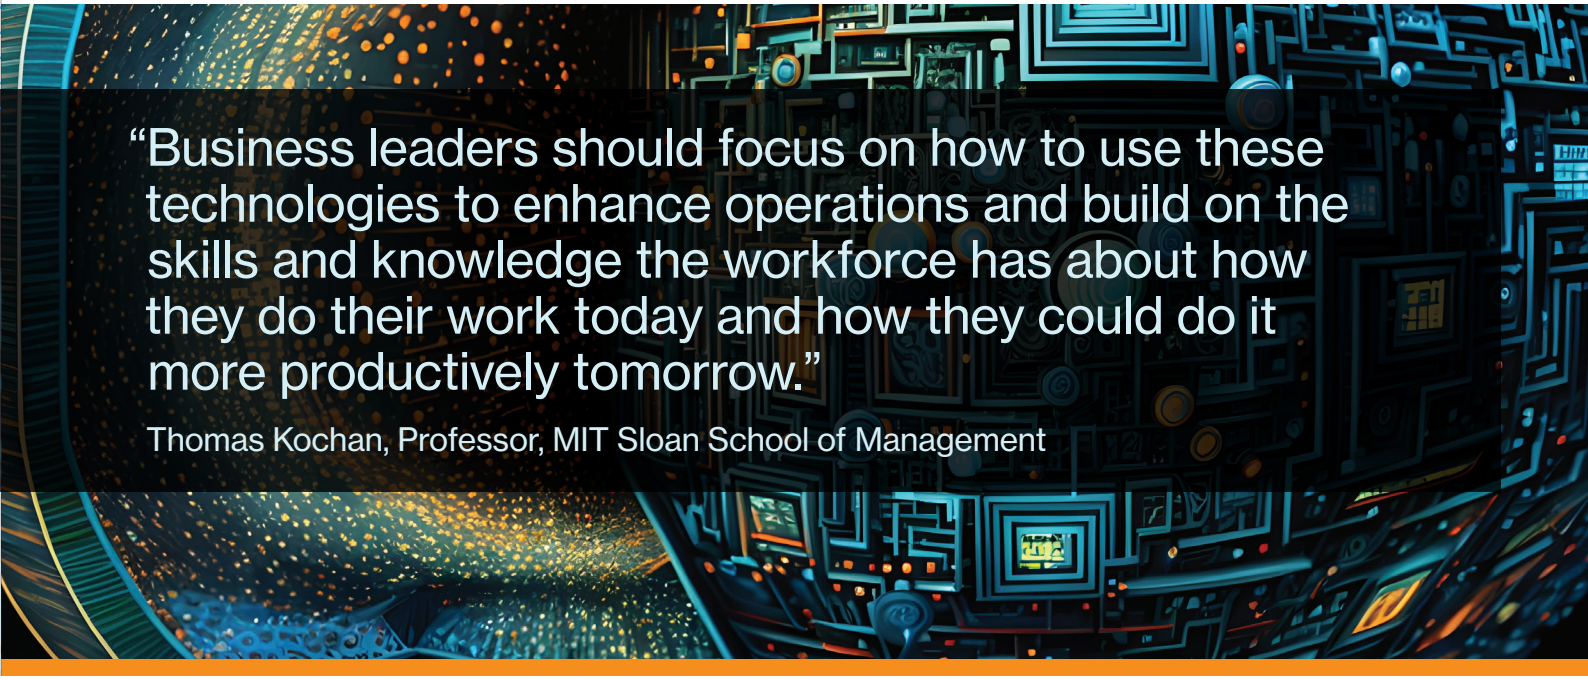

exponentially thereafter. “With robotic process automation, you saw firms would start with one robot on one task, then it would be 10 by year two, 100 in year three, 1000 in year four,” says Azhar. Analogously, he predicts, “The amount of generative AI experimentation going on will lead to rapid expansion in two or three years’ time.”

• **Partner up to build capabilities for scale.** Very few companies will travel alone in the generative AI era. A thriving support ecosystem, ranging from large tech companies to startups, can help organizations develop capabilities for model tuning, productization, safety, and governance. Neither Big Tech nor startups enjoy an unassailable advantage. Large tech firms are moving quickly to embed generative AI in their product ecosystem, but many executives are looking to specialized startups, especially in sensitive or highly regulated industries. A key risk for these startups is whether step-change improvements in foundation models themselves will encroach on their terrain.

• **Workforce shifts are not predestined: today’s decisions could spread productivity and value.** Dystopian unemployment forecasts may overestimate the power of automation and undervalue the know-how of employees. In sectors with supply-demand mismatches, such as manufacturing, generative AI could be a considerable boon. But AI could, if not co-designed and developed with the workforce, lead to deskilling, increases in cognitive load, and declines in job satisfaction. Executives should look to the potential of technology to enhance customer and employee experience, not merely to cut costs. “Business leaders should focus on how to use these technologies to enhance operations and build on the skills and knowledge the workforce has about how they do their work today and how they could do it more productively tomorrow,” says Kochan. “Then we can produce more joint gains for the enterprise, for the workforce, and society.” Future of work committees that involve a cross-section of the workforce and collaborations with labor organizations could foster a more cohesive strategy.

• **New laws and institutions are warranted to manage AI risk, but these will take time to establish.** Regulators are right to consider new and updated laws and legislation for AI, but these should be specific, and informed by a broad stakeholder community. In the meantime, mounting legal clashes over IP and copyright could subside as industries and tech companies work together on new licensing and revenue models. In what may be AI’s “Napster moment,” the current period could give rise to new business models and consumer norms. This may also be an opportunity to fundamentally revisit some assumptions and norms in areas like copyright.

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 thought leaders worldwide for surveys and in-depth interviews.

Footnotes

1. Ashish Vaswani et al., "Attention Is All You Need," *Advances in neural information processing systems* 30, 2017, https://arxiv.org/abs/1706.03762.
2. Gary Marcus, "What If Generative AI Turned Out to Be a Dud?" Marcus on AI, August 20, 2023, https://garymarcus.substack.com/p/what-if-generative-ai-turned-out.
3. Anissa Gardizy and Kevin McLaughlin, "How AWS Stumbled in AI, Giving Microsoft an Opening," *The Information*, August 30, 2023, https://www.theinformation.com/articles/how-aws-stumbled-in-ai-giving-microsoft-an-opening.
4. "New PitchBook Survey Reveals Venture Investors Bullish on AI and Cautious on Fundraising," *PitchBook*, May 1, 2023, https://pitchbook.com/media/press-releases/new-pitchbook-survey-reveals-venture-investors-bullish-on-ai-and-cautious-on-fundraising.
5. Jacob Robbins, "Generative AI startups jockey for VC dollars," *PitchBook*, April 14, 2023, https://pitchbook.com/news/articles/Amazon-Bedrock-generative-ai-q1-2023-vc-deals.
6. "Exploring opportunities in the generative AI value chain," McKinsey & Company, April 26, 2023, https://www.mckinsey.com/capabilities/quantumblack/our-insights/exploring-opportunities-in-the-generative-ai-value-chain.
7. "27% of jobs at high risk from AI revolution, says OECD," *Reuters*, July 11, 2023, https://www.reuters.com/technology/27-jobs-high-risk-ai-revolution-says-oecd-2023-07-11/.
8. Imane El Atillah, "Highly skilled occupations such as medicine, law and finance 'may be at risk of automation from AI'," *Euronews*, July 13, 2023, https://www.euronews.com/next/2023/07/13/highly-skilled-occupations-such-as-medicine-law-and-finance-may-be-at-risk-of-ai-automatio.
9. "A new era of generative AI for everyone," Accenture, 2023, https://www.accenture.com/content/dam/accenture/final/accenture-com/document/Accenture-A-New-Era-of-Generative-AI-for-Everyone.pdf.
10. Anna Lena Hunkenschroer and Christoph Luetge, "Ethics of AI-Enabled Recruiting and Selection: A Review and Research Agenda," *Journal of Business Ethics* 178, February 8, 2022, https://doi.org/10.1007/s10551-022-05049-6.
11. Patrick Coffee, "AI Set to Upend the Talent Search at Ad Agencies and Brands," *The Wall Street Journal*, August 7, 2023, https://www.wsj.com/articles/ai-set-to-upend-the-talent-search-at-ad-agencies-and-brands-cf9df34f.
12. J. Scott Marcus, "Adapting the European Union AI Act to deal with generative artificial intelligence," Bruegel, July 19, 2023, https://www.bruegel.org/analysis/adapting-european-union-ai-act-deal-generative-artificial-intelligence.
13. Javier Espinoza, "European companies sound alarm over draft AI law," *Financial Times*, June 30, 2023, https://www.ft.com/content/9b72a5f4-a6d8-41aa-95b8-c75f0bc92465.
14. Matt Sheehan, "China's AI Regulations and How They Get Made," Carnegie Endowment for International Peace, July 10, 2023, https://carnegieendowment.org/2023/07/10/china-s-ai-regulations-and-how-they-get-made-pub-90117.
15. Matthew Hutson, "Rules to keep AI in check: nations carve different paths for tech regulation," *Nature*, August 8, 2023, https://www.nature.com/articles/d41586-023-02491-y.
16. "A pro-innovation approach to AI regulation," Department for Science, Innovation & Technology, March 2023, https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach.
17. "Japan revises IP promotion blueprint policy on AI," June 10, 2023, *Digital Watch*, https://dig.watch/updates/japan-revises-ip-promotion-blueprint-policy-on-ai.

About the partner

Adobe is changing the world through digital experiences. For more information, visit **www.adobe.com**.

About the partner

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets. Enabled by data and technology, diverse EY teams in over 150 countries provide trust through assurance and help clients grow, transform and operate. Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit **ey.com**.

About the partner

Owkin is a TechBio company that combines the best of human and artificial intelligence to ensure every patient gets the right treatment. By understanding complex biology through AI, we identify new treatments, de-risk and accelerate clinical trials, and develop AI diagnostics. Owkin uses privacy-enhancing federation to access up-to-date multimodal patient data that unlocks AI's potential to power precision medicine. We merge wet lab experiments with advanced AI techniques to create a powerful feedback loop for accelerated discovery and innovation. Owkin founded **MOSAIC**, the world's largest multi-omics atlas for cancer research. Owkin has raised over $300 million through investments from leading biopharma companies (Sanofi and BMS) and venture funds (Fidelity, GV and BPI, among others).

Illustrations
Cover art by Adobe Stock.

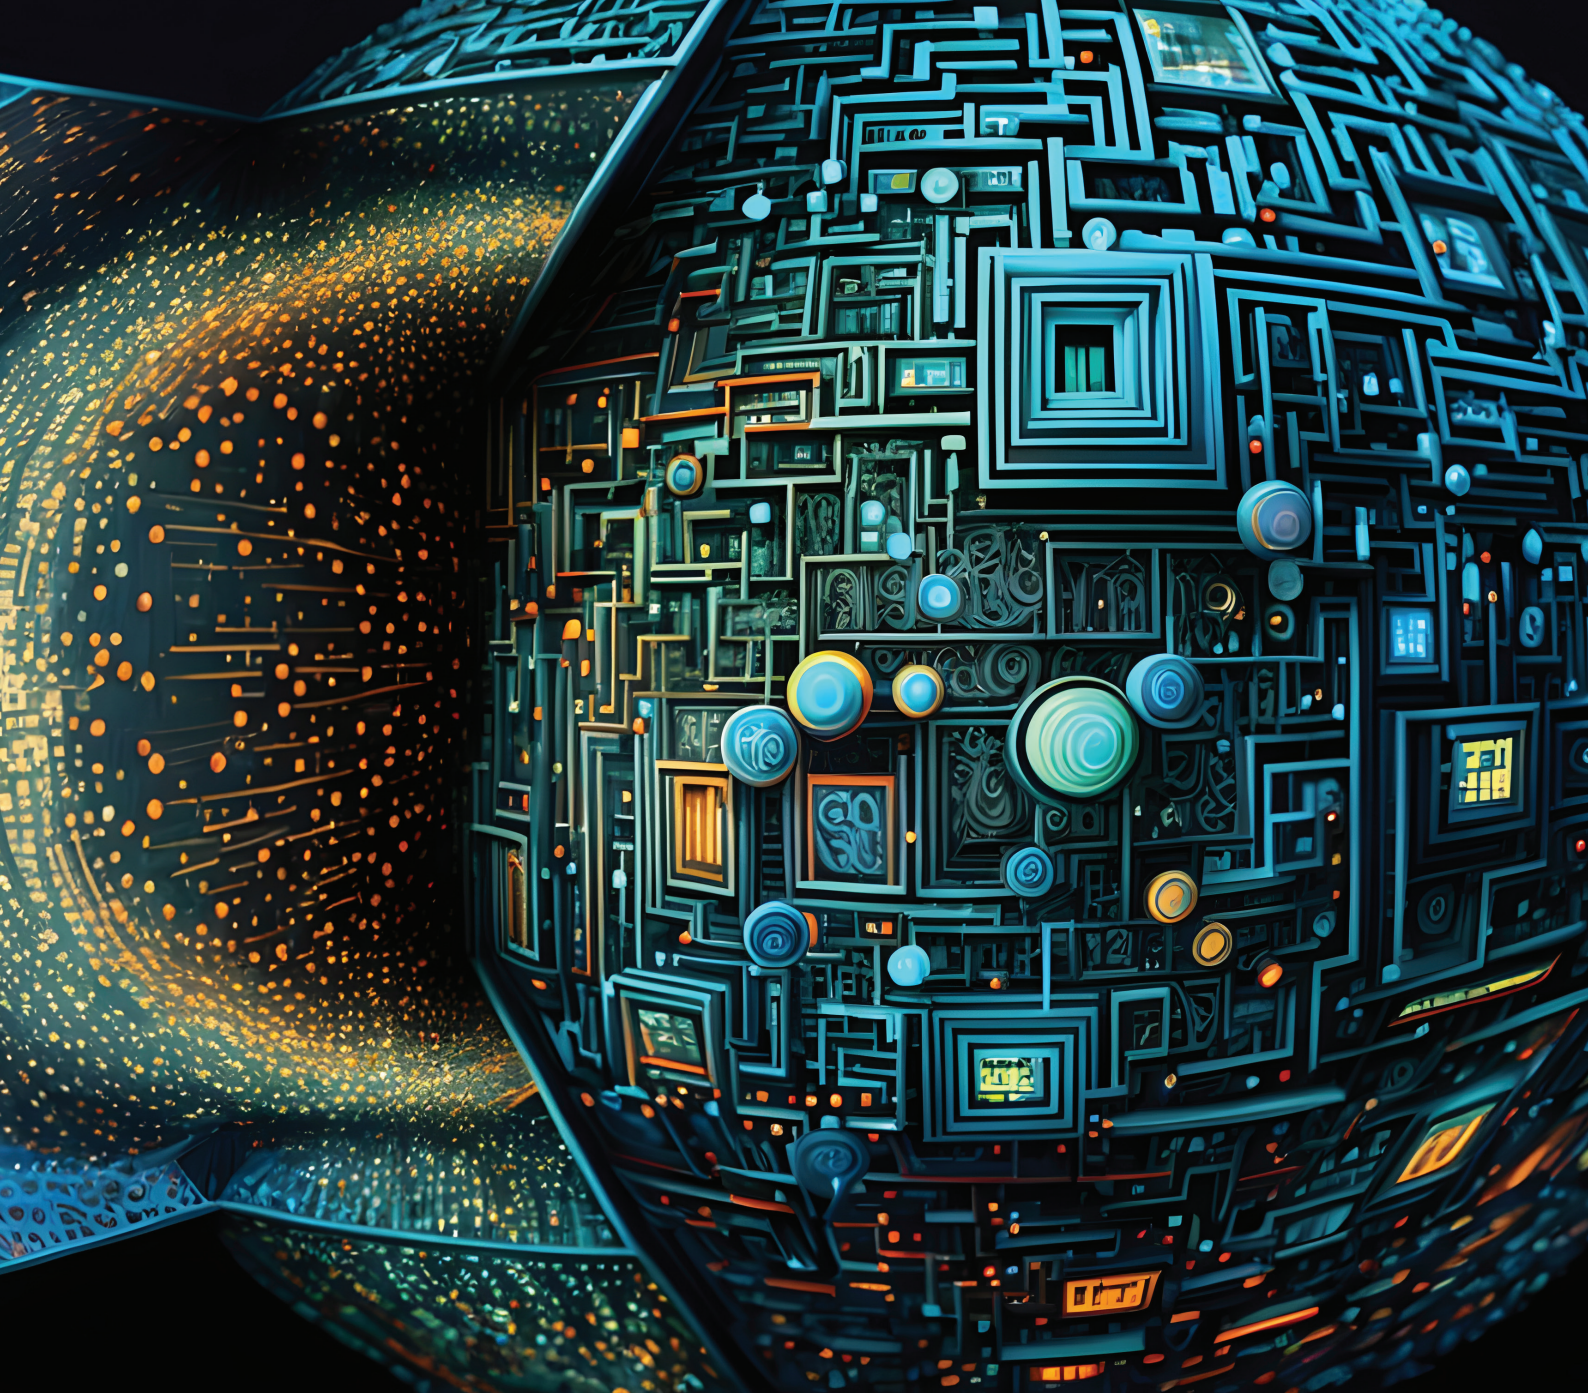

MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com