September 2023

The power of AI and generative AI: what boards should know

AI and generative AI offer new strategic opportunities — and risks, both operational and reputational, that need to be managed. The board has a role to play in overseeing trusted AI systems to help safeguard the company while delivering value.

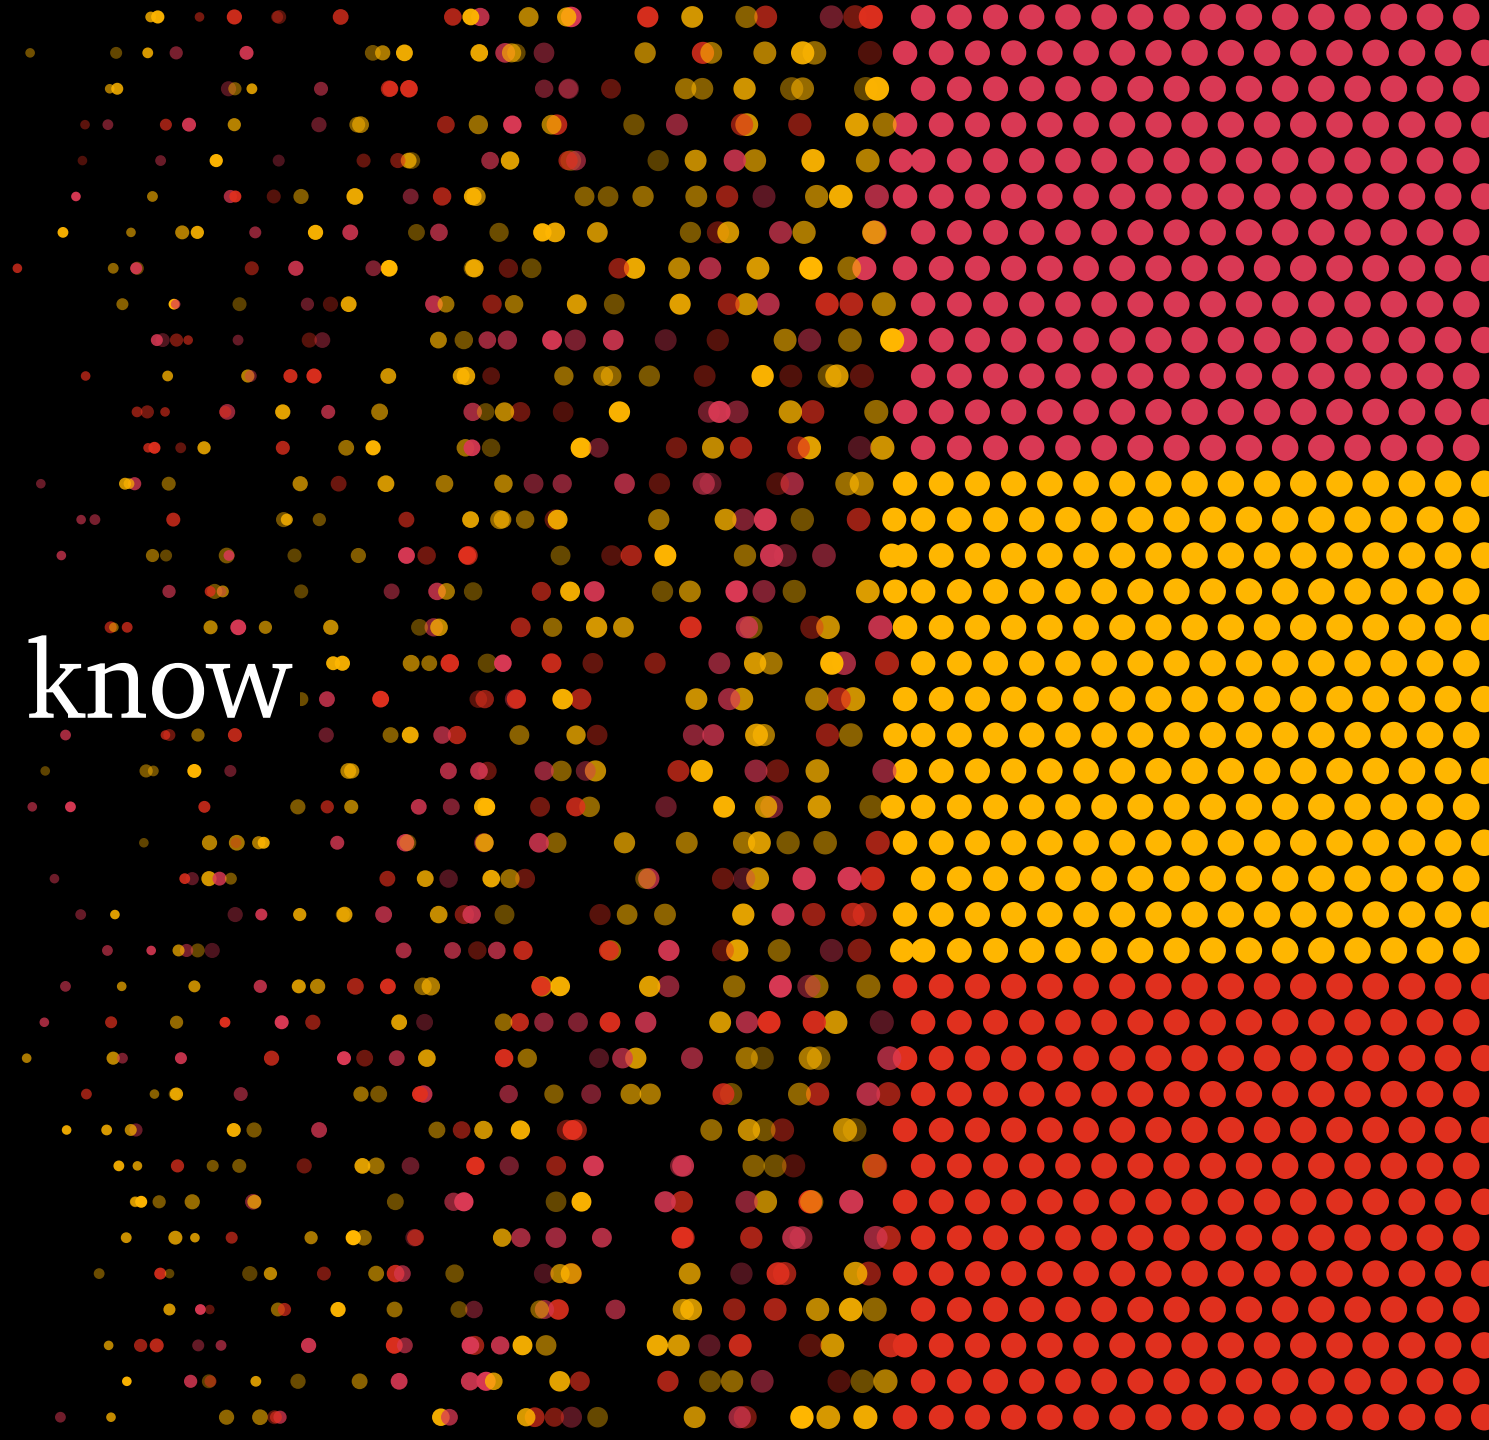

pwc

Table of contents

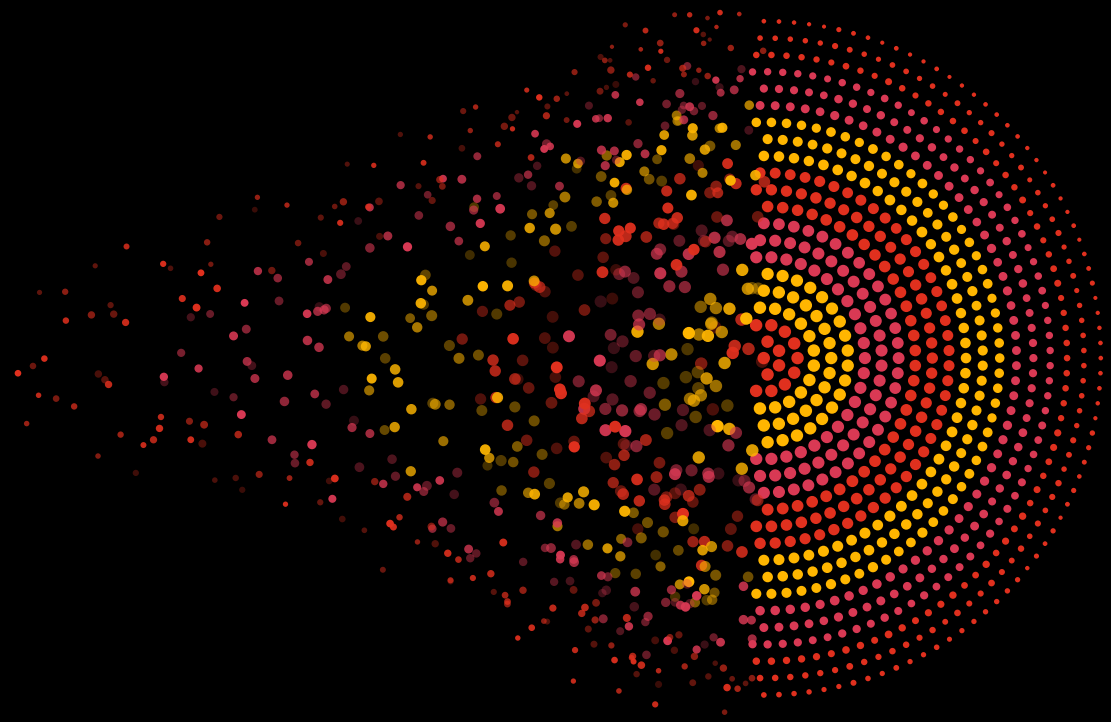

A watershed moment

Artificial intelligence is back in the headlines. The technology has been around in various forms for decades, and in routine use in companies for nearly as long, but the convergence of cloud computing, high-speed computer processing and ubiquitous data have all dramatically increased accessibility and use. “Generative AI,” including ChatGPT, Bard and DALL-E 2, has come to dominate conversations everywhere from social media to strategic discussions. C-suite executives and directors are looking to catch up on developments in AI, with companies discovering new, advanced applications nearly every day.

Long before OpenAI launched ChatGPT into the world, AI moved from a science fiction concept to become part of our everyday lives at home and work, in ways that don’t necessarily call attention to themselves — for example, in shopping suggestions, ride-hailing services, language translation apps, facial recognition technology and voice-activated home devices. At the office, AI applications have become intrinsic to an increasing number of business functions, helping employees efficiently process information and speed the pace of work.

Because most AI applications work in the background and augment existing functionality — think of a biometric smartphone app or a new car’s navigation system — we may not even remember our first personal encounter with AI. But the fanfare surrounding ChatGPT has rapidly raised awareness of the full range of AI applications, reflecting the technology’s transition of AI from the realm of researchers, theorists and data scientists to everyone. It also constitutes a watershed moment for leaders who recognize that business transformation is certain and necessary for their companies to remain competitive, relevant and able to deliver on their core business. Key questions include how, when and where to use AI, the extent to which it may or will change the way the business operates, and how to establish and maintain trust in its use.

To provide appropriate guidance to your companies, boards need to understand this technology’s potential — and its limitations.

Here are four key areas that boards should address for AI and generative AI:

- **Develop a board approach**
- **Understand the strategic opportunities**
- **Oversee risks and controls for trusted AI**
- **Keep up with emerging regulations**

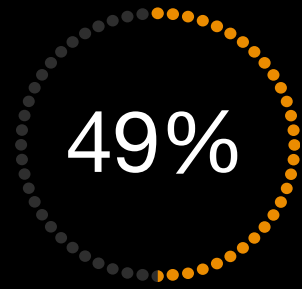

of CEOs think that technology disruptors, such as **AI**, will **impact profitability** to a large extent over the next 10 years.

Source: PwC 26th Annual Global CEO Survey, January 16, 2023: base of 4,410

The evolution of AI…to date

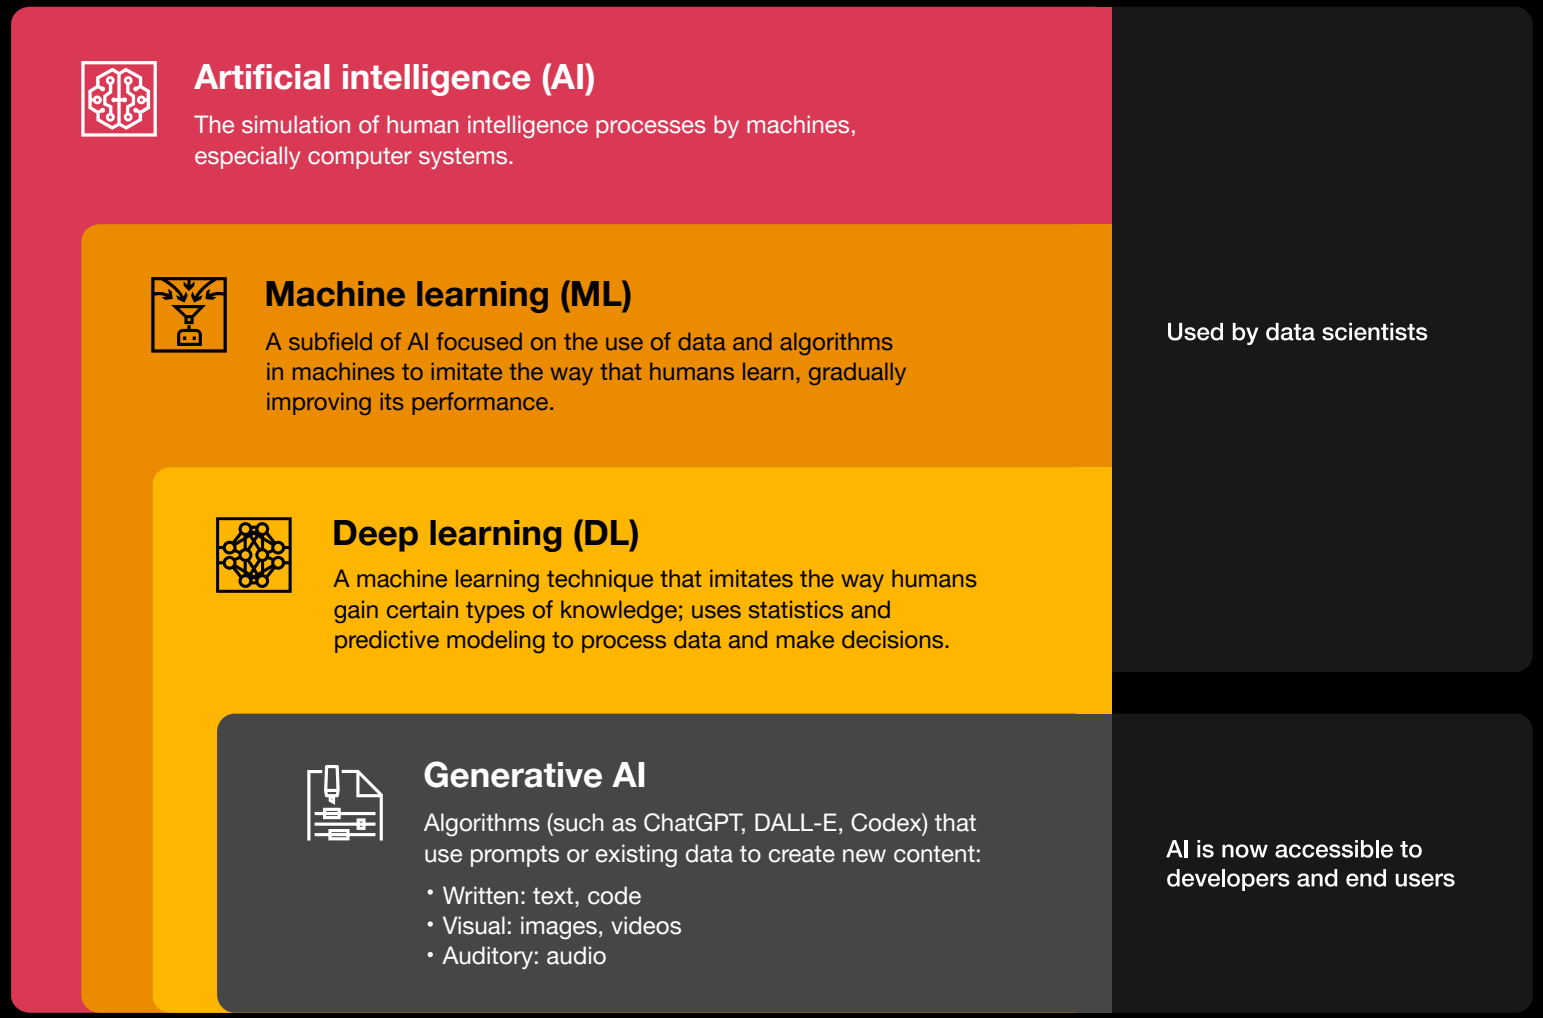

Defining AI and generative AI

In fact, AI is such a familiar yet routinely misunderstood concept that it's worth a refresher: Exactly what is it, how does it work today and how much further does generative AI take the technology?

AI, in our broad definition, is a collective term for computer systems that can sense their environment, think, learn and take action in response to what they're sensing and their objectives. AI enables computers, apps and an ever-growing range of programs and devices to perceive the digital or physical world, quickly process what they perceive, and make decisions, recommendations, and/or take actions that traditionally have required human intelligence. They do this by being designed to recognize and interpret digitized text, sound and images (i.e., the data), with automated analysis (i.e., the algorithm) to help answer questions, identify categories or patterns, suggest solutions, and diagnose or fix problems. As well-designed AI systems take in more and more data, and are maintained to take advantage of that data, the performance of well-designed systems is likely to improve over time and in many cases improve with each use.

Generative AI is so called because it goes one step further. It is a type of deep learning that can create content. Since it often works on plain language commands, it can be remarkably easy to use. Simply put, it can not only process and analyze data but generate fresh prose, images, video, software code and audio. Models are trained on large amounts of data from a wide range of sources. When given user prompts, they can summarize text, generate and refine code, answer questions, engage in a dialogue, create synthetic data for testing and more. And ChatGPT is only one generative AI tool: Solution providers offer many others, in addition to those available in the open-source community. Companies' proprietary generative AI systems, built on available models that integrate a company's unique data and intellectual property in a secure way, can be fine-tuned, engineered and trained to produce content relevant to a company, specific sectors, businesses and functions.

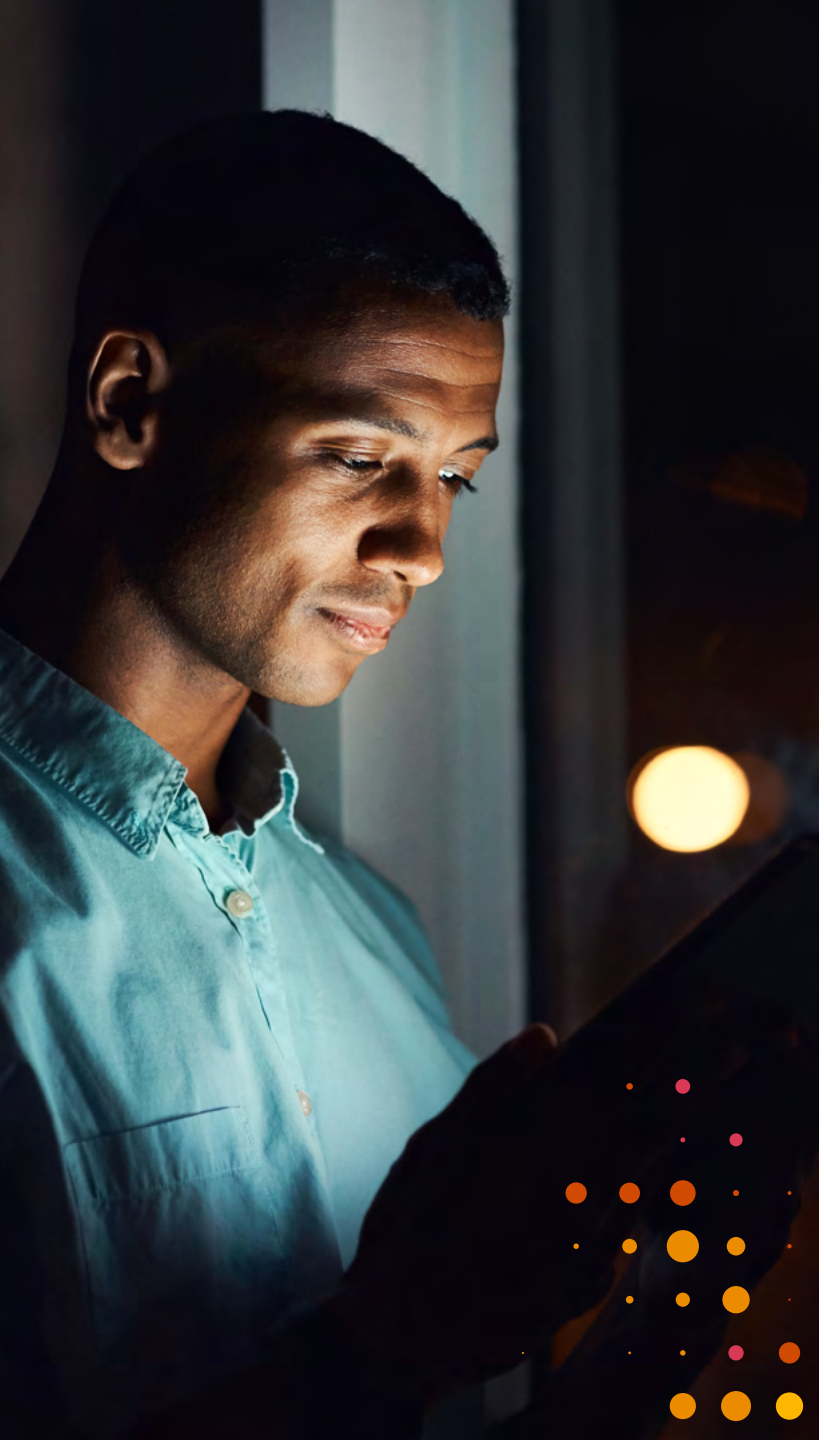

Develop a board approach

Most companies will likely continue to expand their use of AI and explore the possibilities of generative AI. The board's role is to oversee management and focus on how these technologies may impact the corporate strategy and how risks — especially those that are mission-critical or threaten reputational risk — are managed across the company while not limiting innovation.

Boards, in their oversight role, will want to: understand strategic opportunities, oversee risks and controls for trusted AI, and keep up with emerging regulations, as outlined in the following sections.

In addition to these areas, boards will want to think about their approach, specifically:

- **Get educated on AI and generative AI**
- **Review the costs and benefits of this technology**
- **Have a governance model with accountability**
- **Oversee a plan to measure success**
- **Consider communications with stakeholders**

Get educated on AI and generative AI. The first place to start is for boards to increase directors' knowledge of AI and generative AI, tapping both management and external resources to stay apprised of the technologies' growing capabilities, keep up with new use cases and how business models are changing, and risks and responsible use.

When addressing AI and generative AI, directors need to think about the technology and its use from a business perspective. As with the many other areas that the board oversees, its role is to ask management good questions — some examples to start with are shared in this report — and challenge management when appropriate. As the board addresses AI and generative AI, the board can consider whether additional skill sets are needed, or whether to rely on management's skills or those of third parties.

AI could contribute up to $15.7 trillion to the global economy by 2030. Who will get the biggest share of this prize? Those who take the lead now.

Source: PwC's Global Artificial Intelligence Study: Exploiting the AI Revolution

Review the costs and benefits of this technology. Boards should discuss with management the overall benefits and costs to the company. For benefits, it's a matter of reimagining how you get things done — how employees work, how customers engage, and what you sell and how you compete. Costs for employees across the company to be trained and upskilled on AI and generative

AI use may be needed, as management looks to build a culture that encourages employees to learn new skills — and one that incentivizes speed and innovation to capture new opportunities. Additional costs to customize AI to the company and enhance computing power and technology platforms may be needed to run AI and generative AI. Other costs can be to purchase and implement third-party software, establishment of an AI innovation center to develop and test use cases, and others can add up with implementation across the organization if done at scale.

Culture may determine who wins and loses. If your people aren't willing to learn new skills, your organization may not be able to pivot to capture new opportunities.

Have a governance model with accountability. A governance model that drives accountability is critical, beginning with understanding who owns the responsibility for AI governance within the company. An effective governance model enables an organization to evaluate the unique benefit and risk trade-offs associated with its particular technology and individual use cases.

An enterprise-wide management-level AI committee to brainstorm and discuss business opportunities, risks, ethics and set policies can be valuable. Establishing such a committee can facilitate sharing of data, insights and leading practices and help align outcomes with business priorities, furthering organizational buy-in and projects that efficiently deliver a real and sustainable impact.

The significant impact of AI and generative AI on business means that the CEO and C-suite executives have a role to play and participate on the management-level AI committee. The CEO sets the strategy, with special attention to public policy developments, corporate purpose and values. Chief risk and compliance officers are in charge of controls, including governance, compliance and risk management. Chief information and information security officers take the lead on responsible practices, such as cybersecurity, privacy and performance. Data scientists and business domain specialists identify the appropriate AI solution for the use case and apply responsible core practices during development. They also formulate problems and prompts and validate and monitor outputs.

Oversee a plan to measure success. As companies decide their path forward with AI and the business use cases that management is prioritizing, there may be a large digital transformation across the company and a substantial investment required. When a big investment is being made to transform the company, the board will want to understand the digital transformation strategy and plan around AI, and how it aligns with the business strategy. It should oversee the transformation and major investment by understanding management's goals, milestones and impact on talent strategy, infrastructure and other areas for success. Another area to focus on is whether management is taking a holistic rather than piecemeal approach to AI, to reap the advantages of using the technology at scale.

Responsible AI practices can help your stakeholders trust the outputs that your AI produces. This can offer you a competitive advantage in both the market and the competition for talent.

Consider communications with stakeholders. The board should also examine how the company communicates its AI story both internally and externally to stakeholders, focusing on the strategic changes management is making to remain relevant and competitive in today's rapidly changing business environment and the safeguards employed to protect staff, customers, the organization and other stakeholders.

The board's focus: develop a board approach

- How is the board continuing to get educated on AI and generative AI to better understand the technology?
- How have the costs and benefits of AI use been evaluated?
- Who at the management-level is accountable for AI? Is there a management-level committee focused on AI to address opportunities, risks, controls and establish policies?
- For major investments in AI and generative AI, what is management's strategic plan, and does the board have an oversight process in place to monitor the investment's success?
- Have we considered communications to stakeholders about our AI use? What are we communicating to investors?
- Given the pace of change, how do we enable agility to continue to innovate while addressing risks?

Understand strategic opportunities

AI offers businesses real benefits for and solutions to a wide range of problems at all levels of a company. By identifying and analyzing patterns and trends in oceans of data that no person or team could possibly make sense of by themselves, AI can facilitate decision-making, perform repetitive tasks and drive faster business solutions. Companies are using AI-powered simulations on complex subjects to support risk management activities such as supply chain disruptions. And the technology's use cases, in most if not all sectors, are only growing. Intelligent retail spaces recognize shoppers through customer profiles and loyalty programs and provide a customized shopping experience, adapting displays and unlocking promotions. The healthcare industry is leveraging AI for earlier disease detection; financial firms use the technology to assess potential borrowers' creditworthiness; automakers base self-driving capabilities on AI.

The rise of generative AI, expanding the technology's capabilities, has set off an effort to identify ambitious new business use cases, from code development to information extraction and summarization and far beyond. Generative AI has the potential to automate and enhance aspects of almost all business operations and boost productivity, from customer service to software development to data analytics. For example, it can improve customer engagement by drawing on databases to make self-service both more conversational and more powerful, automate high-volume tasks such as processing insurance claims and communications or software development tasks, and make it easier for teams to understand all the unstructured data that matters, from contracts and invoices to customer feedback, policies and performance reviews.

Use cases such as these and many others could impact marketing departments, core business operations and every other function, leading to new product or service development and new ways of working. As the demand for this technology continues to grow, so do its capabilities. While not all use cases identified are feasible today, components may still be; use cases can be improved upon as the technology improves. As AI language systems continue to advance in sophistication and use, innovation will likely only accelerate.

Generative AI business use cases

Generative AI can...	Example capability	Use cases
Create. Query. Transform. Summarize. Q&A.	**Text** →	• Contact center dialog • Improve marketing content • Produce product documentation • Test plan creation • Summarize service requests • Analyze customer feedback • Text transformation (e.g., translation, style, personalization)
	Code →	• Generate python code • Improve and optimize structured query language • Explain and document code • Update code to new standards and platforms
	Data →	• Validate data • Generate sample and test data • Understand differences in data • Fill in missing data points

Even companies with well-integrated AI applications across the organization will need to work out what value generative AI will offer, including determining the use cases that can help the company address disruptive opportunities, business transformation, competitive pressures and operational pain points. It's important to prioritize business use cases based on potential value versus risk and to create a secure space for developing and testing generative AI applications. And while speed is critical, not every company can or should be an early adopter. After all, new opportunities are certain to continually materialize as generative AI evolves along with notable industry- or task-specific applications.

Putting together an effective generative AI strategy shouldn't be too lengthy of a process. Think in terms of weeks, not months, knowing it will continue to evolve.

As companies think through their strategic priorities, they will likely need to engage and work with leading AI technology providers; most AI systems, established and new, come to the business via third-party applications that are publicly available, secured through a vendor relationship or potentially customized on top of a cloud platform. Boards will want to understand who the company plans to work with and the potential related risks, particularly for newer vendors who will now have access to the company's sensitive data.

The rapid onset of generative AI via large language models (LLMs) has resulted in numerous LLM vendors entering the market as companies seek to develop customized solutions to meet their evolving business needs. Companies are using their own data and content to adapt "out-of-the-box" LLMs to develop prioritized use cases. When driven by well-designed prompts using simple user interfaces, called "prompt engineering," an LLM can deliver company-specific, value-added outputs (e.g., draft language, images). Importantly, to help safeguard proprietary information, companies should confirm that policies and processes are in place to address information that is consumed by and shared as outputs from this technology.

AI and generative AI opportunities also will transform the workplace. Knowledgeable employees will be able to do far more, and do it more quickly than today. But that means, companies should upskill talent to learn to use the tools and likely will require new talent acquisition skills that have to be addressed as part of the talent strategy.

52%

of employees globally expect to see some **positive impact of AI** on their career over the next five years

31%

say it'll increase their **productivity/ efficiency at work**

Source: PwC Global Workforce Hopes & Fears Survey, June 20, 2023: base of 54,000

ChatCompanyABC: creating customized use cases

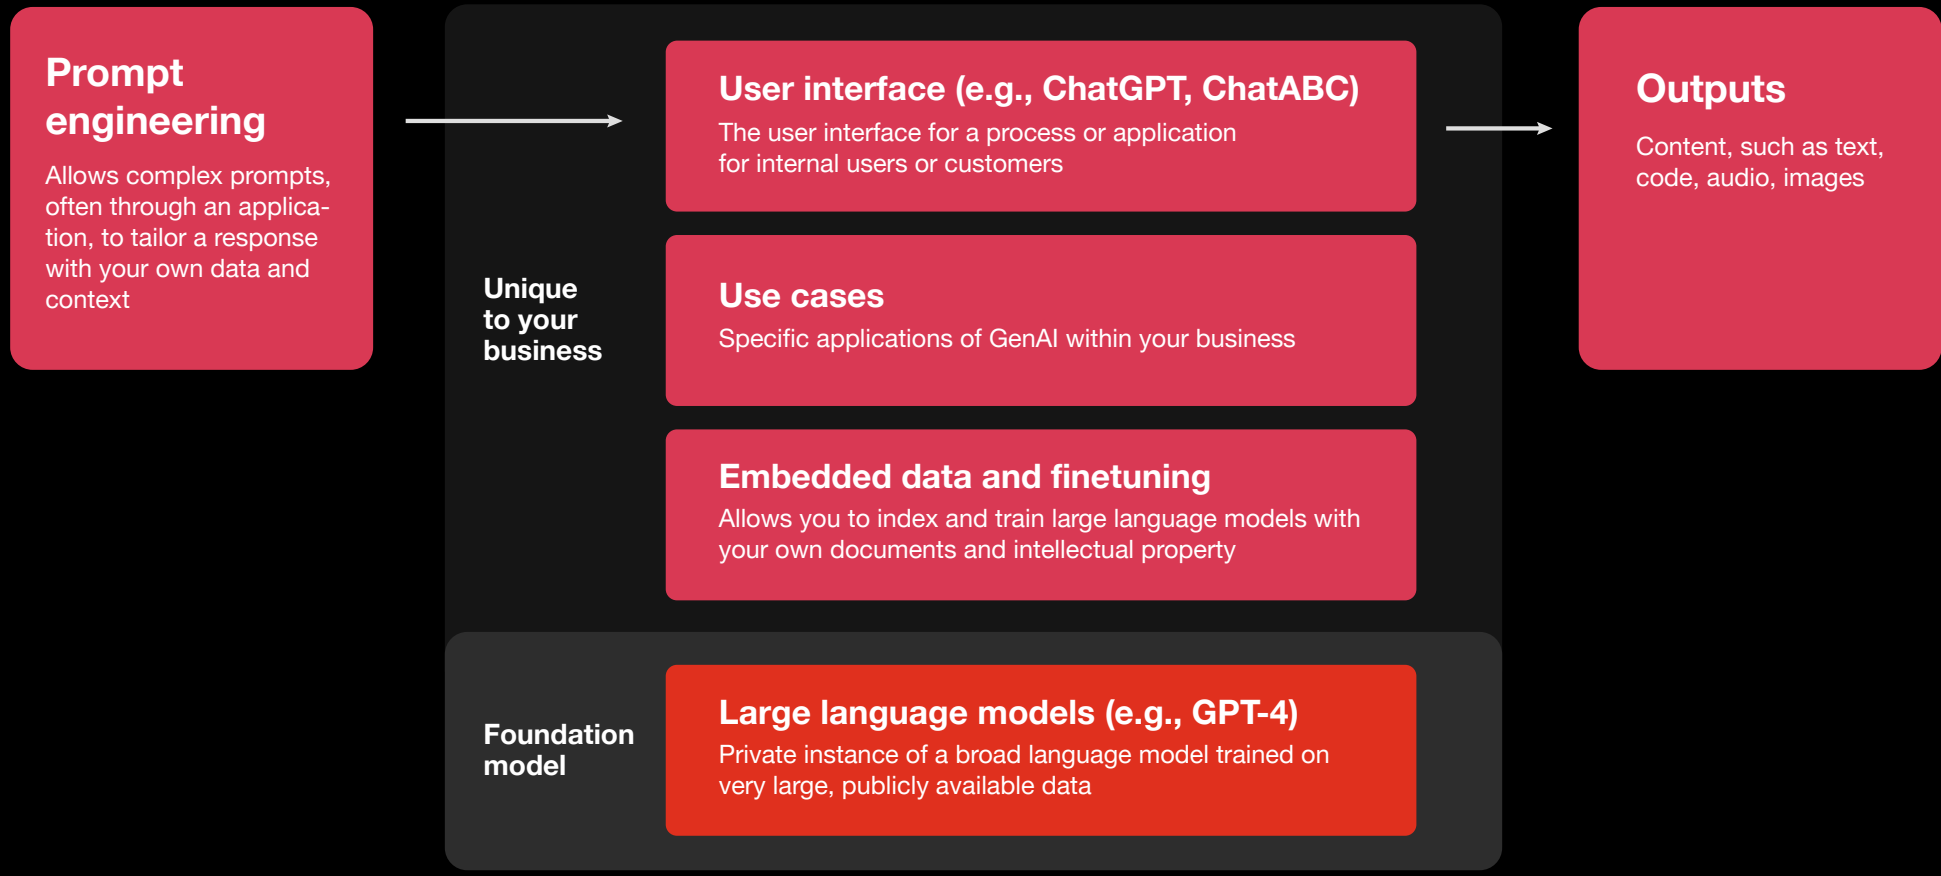

The board's focus: understand strategic business opportunities

- What are the prioritized business use cases for our company? Are we looking at adoption at scale and in a holistic way? Are these use cases complimentary with our other strategic business initiatives?
- How are our competitors starting to use new AI applications to achieve differentiated outcomes?
- Have we considered engaging with an external AI specialist to help validate our prioritized business use cases?
- Are the opportunities to disrupt ourselves (before we are disrupted) within our risk appetite?
- What technology and infrastructure are needed to address AI business use cases? Are we engaging in discussions with a third-party technology provider? What processes do we have in place to evaluate the possible risks related to a third-party provider?
- How will the talent strategy change? How will we engage and upskill existing talent to use the technology and what new skill sets will be needed for success? Has management considered the social implications and long-term talent development implications?

Oversee risks and controls for trusted AI

AI has the potential to create huge economic benefits, but as business applications grow in number and scope, so do risks, both operational and reputational. Generative AI only exacerbates issues, since what makes these systems so powerful is the very thing that concerns so many: Their creators have coupled incredible amounts of public and proprietary content with tremendous computing power, and — in the case of publicly-available LLMs — there is no way for any company to independently validate that all of the data used is accurate, unbiased, and representative or even relevant to a specific business need. Additionally, whereas most AI development was long assigned to IT and data specialists with experience coding and testing these systems, generative AI serves as the latest in a push to "democratize" AI: Now anyone with an internet connection can direct a generative AI application to perform certain tasks or generate certain outputs.

AI-based risks have to be assessed and managed. They include data security and privacy, with readily available applications reducing barriers of entry for threat actors and data being newly vulnerable to not only theft and loss but to perpetuating biases. Companies also face legal and reputational risks, with potential exposure of intellectual property and proprietary information and — if not properly controlled — inaccurate outputs and communications of false data.

The board should set the tone about responsible AI use and make sure that directors and management are clearly aligned as to its responsible use. Directors will want to understand how management has implemented proper policies, controls and processes to reduce possible bias and error in algorithms and data, working to confirm that systems cannot be manipulated — either inadvertently or deliberately — to generate incorrect outcomes. The goal is to develop and design AI systems' practices and outcomes that are trusted and aligned with corporate values, stated strategic objectives and regulatory requirements.

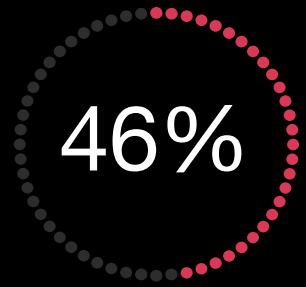

of executives say they're very prepared for an AI algorithm risk or failure

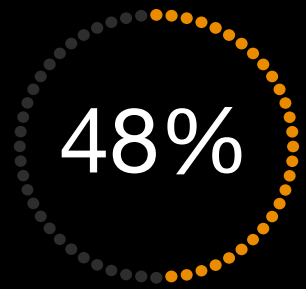

say they're making efforts to reduce the likelihood of such an event a high priority

Source: PwC 2023 Trust Survey, March 2023: base of 500

Characteristics of trusted AI systems

With anxiety rising about the technology's usage, developers, government, and industry organizations and others have created frameworks for boards and companies to help develop trusted AI while capitalizing on its tremendous benefits. The US Department of Commerce's National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF), a voluntary framework, is one of the more frequently adopted.

The AI RMF highlights the importance of establishing responsible AI governance throughout the organization. The framework lays out how organizations can identify, prioritize and manage AI risks. It also describes functions to manage AI risk with governance at the center; and map, measure and manage functions surrounding it.

As the NIST framework suggests, management — with board guidance and oversight — needs to design, develop, deploy and use AI systems as part of a carefully governed, visible system. Managing the multitude of AI risks, particularly those involving generative AI, demands a focus on responsible practices, from keeping processes interpretable and explainable to maintaining data security and fairness. Moreover, as AI-based applications become more powerful in the coming months and years, the burden on companies to capitalize on the benefits and manage the risks of AI is expected to expand dramatically.

If there's a golden rule for responsible AI development, acquisition and deployment, it's this: It's better to implement trust and ethics by design from the start rather than race to close gaps after systems are up and running.

NIST AI Risk Management Framework

Source: NIST AI 100-1 AI RMF 1.0

PwC's approach to Responsible AI

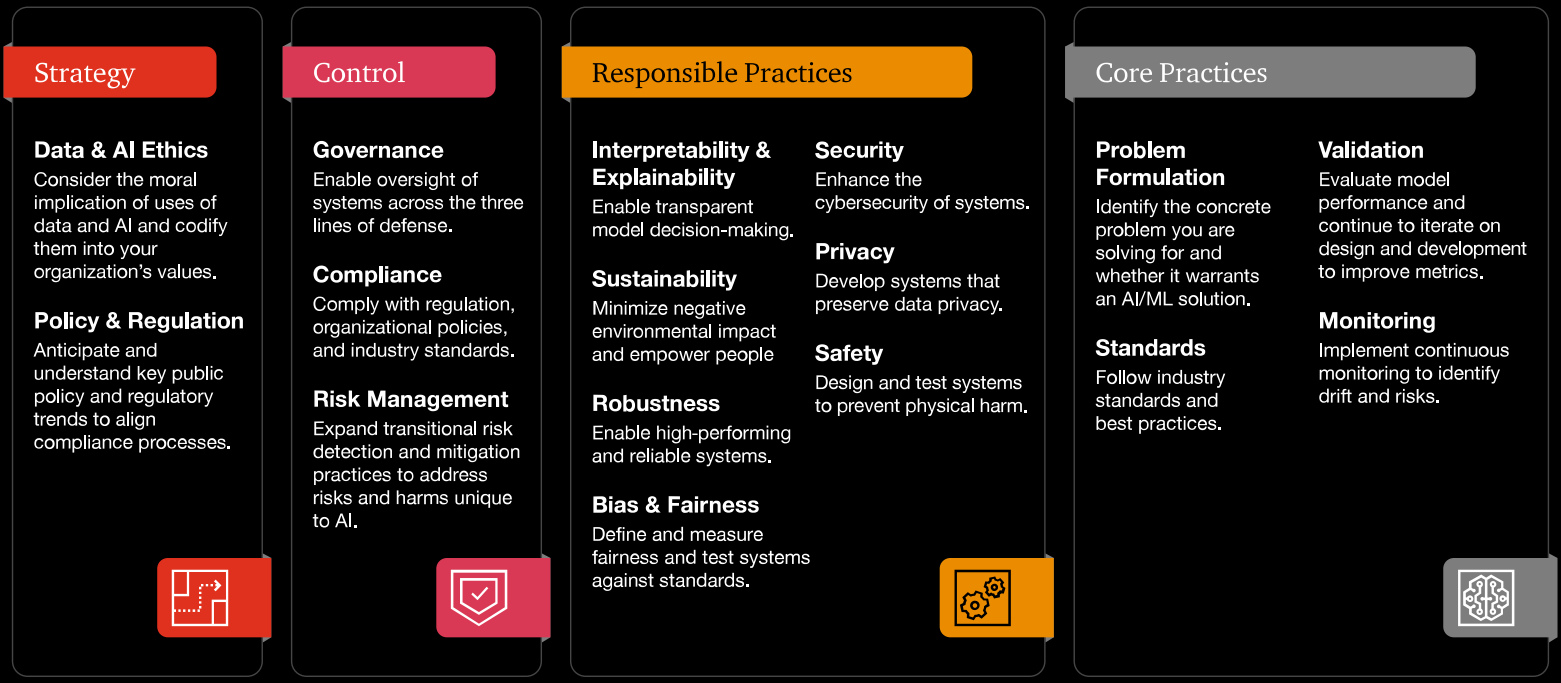

The good news is that — consistent with the NIST AI RMF and all emergent and emerging global policies around the responsible use of AI — PwC's approach to Responsible AI is able to equip companies to identify the risk management dimensions around AI. This toolkit is customizable to use by companies as a whole and system owners alike — regardless of sector, vertical, business model or maturity of their AI systems. Because AI risk will vary even in the same organization — for example, the risks of using AI by COOs and CFOs are different than those of CHROs — the toolkit is able to flexibly serve each CXO based on their respective roles and responsibilities and create a shared responsibility for working together to manage risks when implementing AI systems.

To govern an organization's AI strategy and practice, models should be transparent and thoroughly documented, able to be explained and validated — by humans, not AI — to regulators, shareholders, employees and other stakeholders. One of the central features of AI is that it learns complex relationships from complex data. Often, this data is amassed over time and the more data — especially high-quality, company-specific data — the better the model is able to solve a specific business use case. Nevertheless, confirmation through humans is needed to confirm that the system, fed by relevant, appropriate and reliable data, produces and continues to produce results that are within an expected range of acceptable outcomes. Companies will need to define clearly what the AI models do and how they should perform so that they work the way we expect them to, and boards will want to oversee management continues to monitor the use cases and their outcomes.

Building trust in these systems requires management to establish processes to effectively test and evaluate AI systems. Existing functions, like privacy, IT, data and security governance, risk management and internal audit may be engaged at various points to build this trust. Some companies may even choose independent or third-party evaluation for more critical or more risky applications.

The board's focus: oversee risks and controls for trusted AI

- Has the company considered the responsible use and societal impact of the use of AI technology? How does it align with our corporate culture?
- What framework(s), if any, is the company using to implement AI technology responsibly?
- What policies, controls and processes does the company have in place to safeguard AI and generative AI models against related risks, misuse and unauthorized use? Are they aligned with established legal, privacy, security, and ethics policies and procedures? Is the board setting the tone on responsible AI use?
- What are the processes to independently validate and verify that AI systems are operating in a manner consistent with stated policies and objectives? Does internal audit have a role?
- Does the company's enterprise-wide risk management program incorporate AI-related risks and data protection?

Keep up with emerging regulations

The rapid rise of generative AI has spurred regulators to raise concerns about AI practices, quickly advancing regulators' new rules and initiatives. Daily headlines seem to be filled with new ways to regulate AI from the White House, federal agencies, state legislatures and others around the globe — and it is not expected to slow down anytime soon. Boards should oversee that management keeps up with regulatory change and maintains compliance, particularly with regard to generative AI.

US regulators have different views on how to address AI. Some see existing laws — such as those on data privacy, data protection, anti-discrimination, fair lending and servicing, and hiring — as effectively applicable to AI; other regulators, along with many companies developing generative AI applications, are calling for more policy and regulations. And some regulators and companies are looking to align with voluntary standards such as the NIST AI framework, while a few state and local governments are taking the initiative to propose their own standards.

Regulators around the world are taking their own approaches to standards. The European Union recently adopted the AI Act, a sweeping regulatory framework intended to do for AI what the General Data Protection Regulation (GDPR) did for data, aiming to account for the significant uptick in interest around generative AI and foundation models, subjecting them to more stringent requirements and applying copyright laws to generative AI source content. In China, regulators have moved to ban the use of deep fakes, a category of generative AI used to create images, video and audio of real people, often for the purpose of manipulation.

The board's focus: keep up with emerging regulations

- Have we mapped our planned use of AI applications to existing laws and regulations? Are there any concerns that need to be addressed?
- Have we performed legal due diligence on contracts, including those with third parties, and intellectual property to address the impact of AI and generative AI?
- How are we keeping up with compliance and legal issues that current and prospective AI initiatives might pose to our use of the technology? Do we have sufficient resources to do so?

Conclusion:

AI and generative AI will continue to provide companies with opportunities for greater applications and innovation, and with risk concerns, cost benefit considerations and regulatory requirements to consider. As companies determine how to use this technology responsibly, and build trust in their applications, oversight and guidance will be an important role for boards to play as ever.

The board's focus: questions to ask about AI and generative AI

Develop a board approach

- How is the board continuing to get educated on AI and generative AI to better understand the technology?
- How have the costs and benefits of AI use been evaluated?
- Who at the management-level is accountable for AI? Is there a management-level committee focused on AI to address opportunities, risks, controls and establish policies?
- For major investments in AI and generative AI, what is management's strategic plan, and does the board have an oversight process in place to monitor the investment's success?
- Have we considered communications to stakeholders about our AI use? What are we communicating to investors?
- Given the pace of change, how do we enable agility to continue to innovate while addressing risks?

Understand strategic business opportunities

- What are the prioritized business use cases for our company? Are we looking at adoption at scale and in a holistic way? Are these use cases complimentary with our other strategic business initiatives?
- How are our competitors starting to use new AI applications to achieve differentiated outcomes?
- Have we considered engaging with an external AI specialist to help validate our prioritized business use cases?
- Are the opportunities to disrupt ourselves (before we are disrupted) within our risk appetite?
- What technology and infrastructure are needed to address AI business use cases? Are we engaging in discussions with a third-party technology provider? What processes do we have in place to evaluate the possible risks related to a third-party provider?
- How will the talent strategy change? How will we engage and upskill existing talent to use the technology and what new skill sets will be needed for success? Has management considered the social implications and long-term talent development implications?

Oversee risks and controls for trusted AI

- Has the company considered the responsible use and societal impact of the use of AI technology? How does it align with our corporate culture?
- What framework(s), if any, is the company using to implement AI technology responsibly?
- What policies, controls and processes does the company have in place to safeguard AI and generative AI models against related risks, misuse and unauthorized use? Are they aligned with established legal, privacy, security, and ethics policies and procedures? Is the board setting the tone on responsible AI use?
- What are the processes to independently validate and verify that AI systems are operating in a manner consistent with stated policies and objectives? Does internal audit have a role?
- Does the company's enterprise-wide risk management program incorporate AI-related risks and data protection?

Keep up with emerging regulations

- Have we mapped our planned use of AI applications to existing laws and regulations? Are there any concerns that need to be addressed?
- Have we performed legal due diligence on contracts, including those with third parties, and intellectual property to address the impact of AI and generative AI?
- How are we keeping up with compliance and legal issues that current and prospective AI initiatives might pose to our use of the technology? Do we have sufficient resources to do so?

Resources

For a deeper dive into how businesses can benefit from AI and generative AI, see:

PwC generative AI insights collection

How to leverage generative AI to unlock value and reinvent your business

Managing the risks of generative AI

Guide to scaling generative AI for your business

GenAI: Transform the future of business and lead with trust

What is responsible AI and how can it help harness trusted generative AI?

Unlock the full potential of artificial intelligence at scale — in a way you can trust

For more information:

Maria Castañón Moats
Leader, Governance
Insights Center
maria.castanon.moats@pwc.com

Jennifer Kosar
Leader, Trust and
Transparency Solutions
jennifer.kosar@pwc.com

Barbara Berlin
Managing Director, Governance
Insights Center
barbara.berlin@pwc.com

Tim Persons
Principal, Digital Assurance and
Transparency of US Trust Solutions
tim.persons@pwc.com

Mohini Singh
Director, Governance Insights
Center & Public Policy
mohini.z.singh@pwc.com

Ilana Golbin Blumenfeld
Director, Emerging Technologies
and Responsible AI Lead
ilana.a.golbin@pwc.com