McKinsey & Company

McKinsey Direct

The European Union AI Act: Time to start preparing

A successful digital future depends on responsible use of AI. The EU AI Act marks a significant step in regulating AI systems and could serve as a blueprint for other jurisdictions.

This article is a collaborative effort by Henning Soller with Anselm Ohme, Chris Schmitz, Malin Strandell-Jansson, Timothy Chapman, and Zoe Zwiebelmann, representing views from McKinsey's Risk & Resilience and Digital Practices.

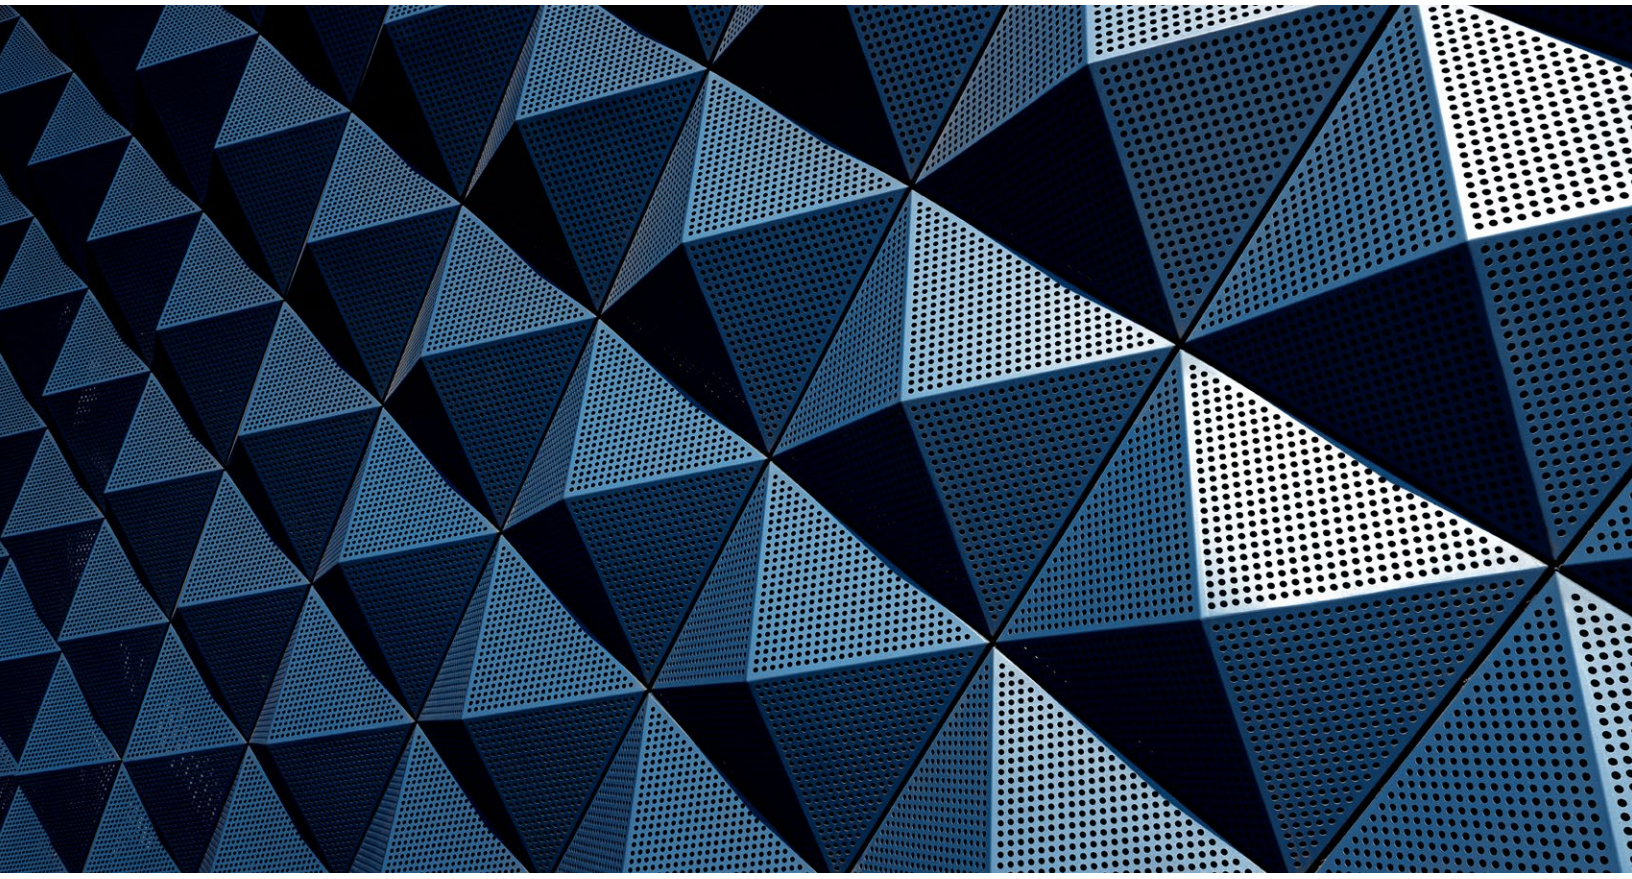

November 2024

Artificial intelligence and generative AI (gen AI) will have a transformative impact on economic growth and productivity. This is especially true for organizations that expect to make changes to their operations using the technology, a recent McKinsey survey shows.[1]

To realize the benefits of AI, organizations need the underlying models and their use to be secure, safe, and trusted. Implementing robust data governance, model-risk, security, and individual-rights management is crucial for responsible AI governance. Together, these pillars create a solid foundation for future digital transformation, and digital trust. According to McKinsey research, trusted organizations have higher margins and better valuations than less-trusted ones.[2] And while only a small contingent of companies are set to deliver this digital trust, organizations that are best positioned to build digital trust are also more likely than others to see annual growth rates of at least 10 percent on their top and bottom lines.

While many organizations embrace these concepts, some still lack fundamental risk controls for the new technologies. In early 2024, McKinsey surveyed 180 EU-based organizations in five sectors about the state of AI governance in the European Union. Seventy-one percent of respondents said their AI risk governance was less than mature, although 65 percent of them said they were already using gen AI (Exhibit 1).

Survey participants expressed concerns in five high-level categories that mirror important considerations for AI: data, model output, security, third-party, and societal risks.

Exhibit 1

Less than 30 percent of survey respondents consider their organization's AI risk governance to have some level of maturity.

Maturity of organization's AI risk governance,[1] % of respondents

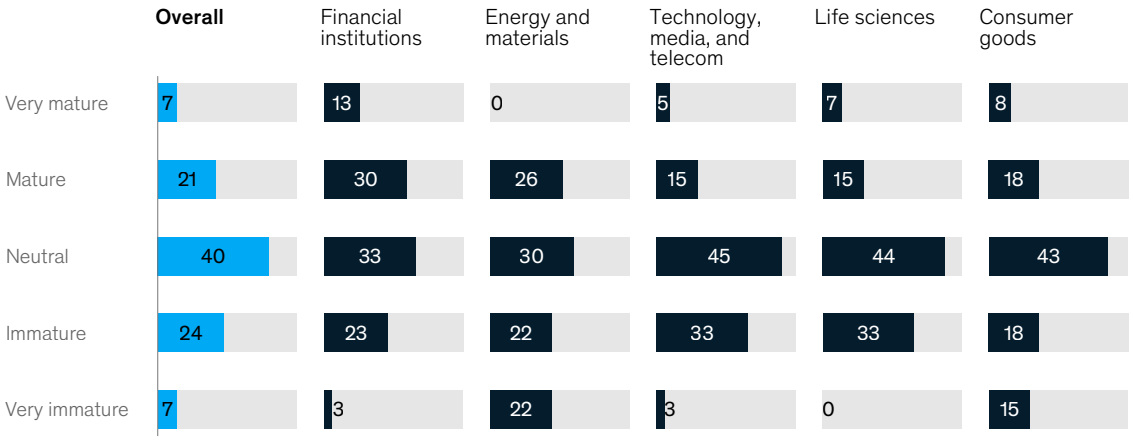

Note: Figures may not sum to 100%, because of rounding.
Question: How mature is your AI risk governance?
Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

[1] "The state of AI in early 2024: Gen AI adoption spikes and starts to generate value," McKinsey, May 30, 2024.
[2] Jim Boehm, Liz Grennan, Alex Singla, and Kate Smaje, "Why digital trust truly matters," McKinsey, September 12, 2022.

Some concerns fall into one category, while others span several. Bias, for example, touches model output, data, and third-party risk. Among the other potential concerns expressed in the survey are discrimination, bad outputs, personal-data leakage, intellectual property misuse, security breaches, and malicious use.

Given everything that could go wrong with AI, standards and policy setters are increasing efforts to control the risks. Regulators globally are introducing regulatory frameworks and guidelines, including in Canada, China, Japan, South Korea, and the United States. The EU AI Act, enacted by the European Union in May 2024, is the world's first general AI regulation to go into effect. Being the first of its kind, the EU AI Act will serve as a test bed for other guidance to follow. In addition, it will have extraterritorial effects because the scope includes AI tools developed in other markets if a tool or its output is applied in the European Union.

Overview of the EU AI Act and its requirements

The EU AI Act aims to "promote human-centric and trustworthy AI while protecting health, safety, and fundamental rights." It will have wide-ranging implications for all affected organizations as the guidance is rolled out over the next two years.

The act sets requirements in four areas: governance, data management, model-risk management, and individual rights. These requirements include risk and quality management, human oversight, AI system documentation and transparency, data management, model-risk governance measures for nondiscrimination and bias, accuracy, robustness, and cybersecurity.

Which requirements apply to each organization depends on two factors: the risk classification and the role of the organization in the AI value chain, which includes providers, importers, distributors, deployers of AI systems, and combinations thereof.

Based on the use case, AI systems are defined as prohibited, high-risk, or non-high-risk. Rules for "prohibited" AI, which includes models that are manipulative or deceptive, are outlined in Article 5 of the act. "High risk" systems are those that could threaten health, safety, and fundamental rights, including those related to critical infrastructure, education or vocational training, employment, access to essential public or private services and benefits (including credit and health insurance), profiling, and law enforcement. "Non high risk" systems, with lower or no regulatory requirements, consist of everything not specifically covered by the other two categories, including AI in video games and customer service chatbots.

Early days of implementation efforts

AI governance and EU AI Act compliance efforts are still in the early days, but organizations already have questions. More than 50 percent of survey respondents said they are not clear on AI act requirements and are unsure of the risk classifications for their AI use cases (Exhibit 2).

Organizations consider themselves most prepared with regard to data management, ahead of governance, model risk management, and individual rights (Exhibit 3).

Even so, data management is still a concern. More than half—57 percent—of respondents said that many data governance requirements remain unaddressed. Specifically, some organizations said there is a lack of clarity in terms of how the General Data Protection Regulation (GDPR) and the EU AI Act will interact.

When asked whether they had already met the act's requirements for the four areas, less than 10 percent of survey respondents said that they had (Exhibit 4).

Exhibit 2

Only 4 percent of survey respondents agreed that the EU AI Act requirements are clear.

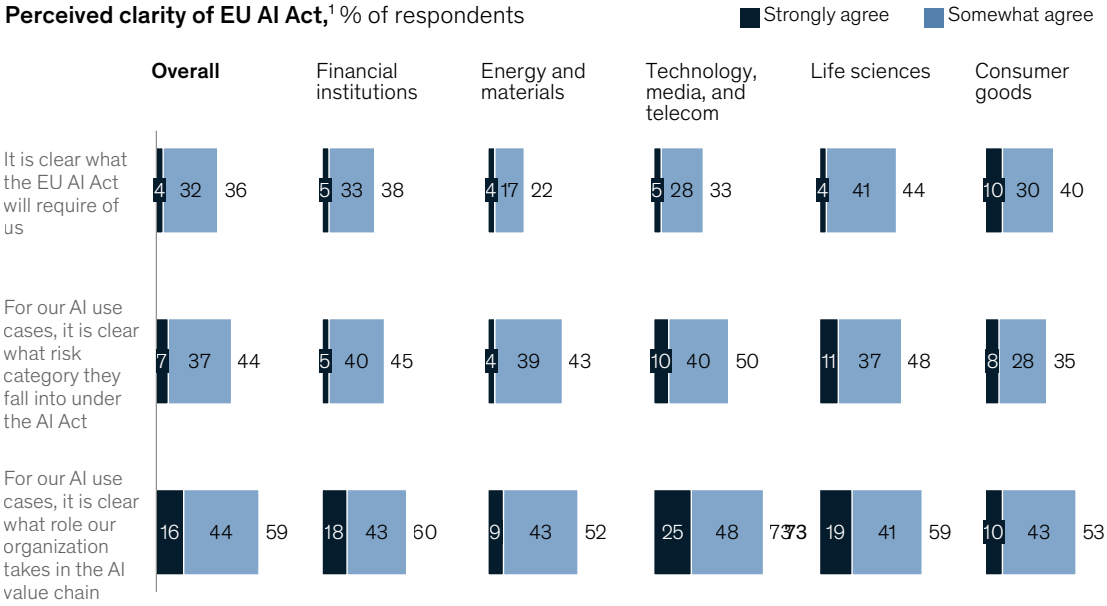

Note: Figures may not sum to totals, because of rounding.
Question: To what extent do you agree with the following statements?
Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

Exhibit 3

Survey respondents consider their organizations somewhat prepared across various dimensions of the EU AI Act.

Self-assessment of EU AI Act governance maturity, averages and ranges

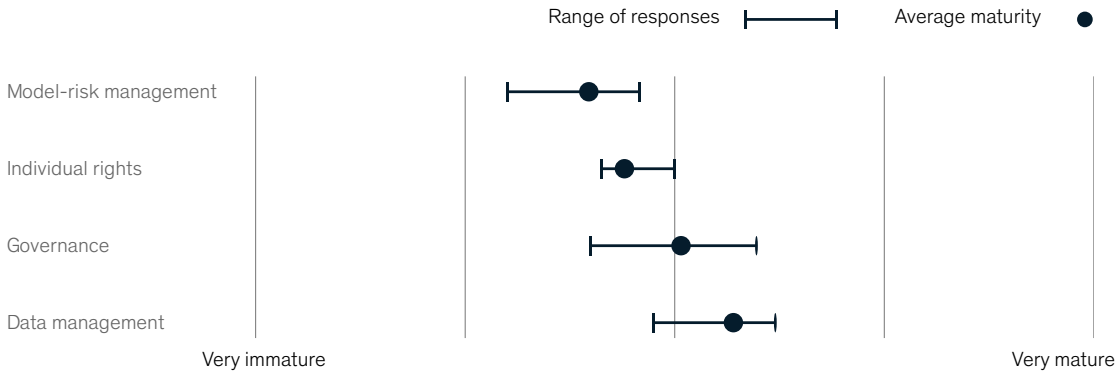

Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

Exhibit 4

Few of the key requirements of the EU AI Act are fully addressed by more than about 10 percent of organizations.

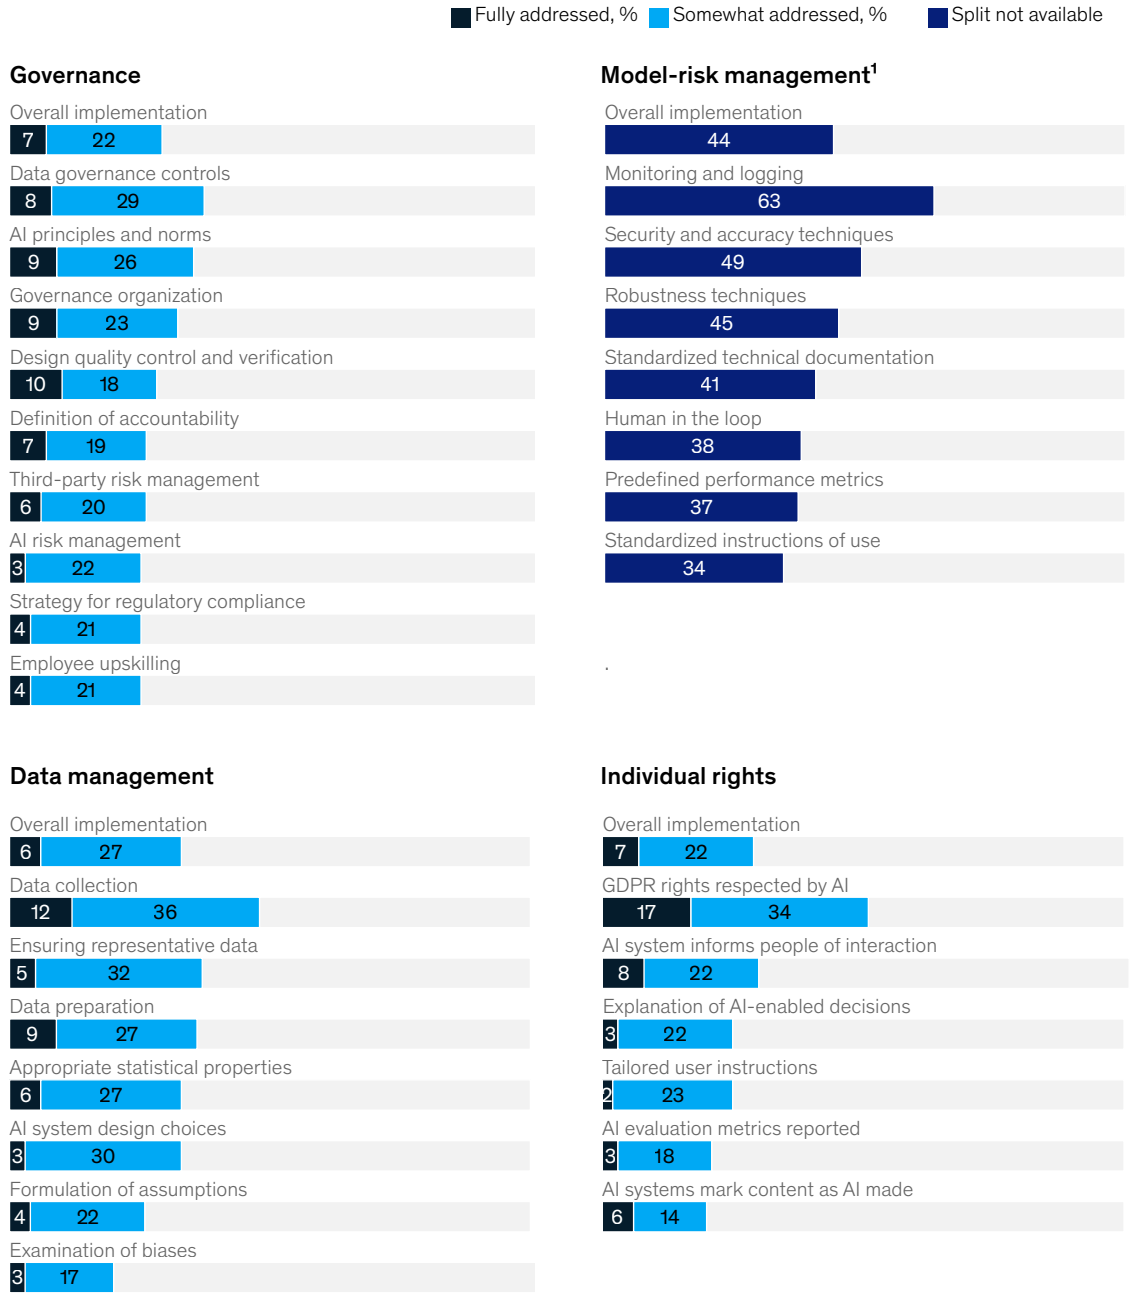

[1]Based on proportion of organizations having technically implemented these measures, not the level at which they have addressed them.
Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

Nearly half of respondents said they had not yet allocated any budget for AI Act implementation, and most that have allocated a budget have set aside €2 million or less (Exhibit 5). There are many reasons organizations aren't spending yet. Some respondents have likely not started responding to AI Act requirements because the rules are so new. Others are focused on aligning their AI remediation efforts to their existing governance structure. Still others are unaware of the upcoming regulatory requirements.

Key challenges facing organizations

Respondents cited a variety of challenges to their efforts to meet the requirements of the AI Act.

Complexity. In some cases, organizations are stalled as they seek clarity and the resources to prepare for complex regulations and technology. Only one in four survey respondents have implemented strategies for regulatory compliance or AI risk management.

Risk governance. About three in ten respondents have developed a mature AI risk governance structure, and only a third said they have a governance organization. Further, about 40 percent lack clear definitions of accountabilities for AI, and about 10 percent say they have fully addressed AI principles and norms.

Encouragingly, nearly half of respondents said they have separate usage guidelines, and more than a third have input and output guardrails in place for external AI models. This likely is a consequence of protecting business-sensitive information and intellectual property as organizations rapidly deployed gen AI tools.

Third-party risk management is also a concern. Less than a third of organizations said they have appropriately addressed AI-related third-party risk.

Exhibit 5

Close to 50 percent of organizations have not yet allocated resources for EU AI Act implementation efforts.

Amount budgeted for EU AI Act implementation efforts,[1] % of respondents

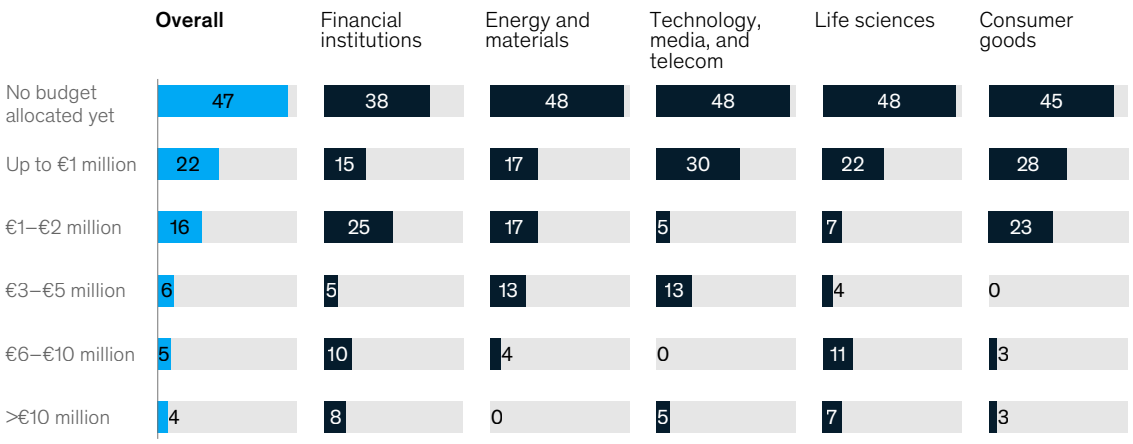

Note: Figures may not sum to 100%, because of rounding.
Question: How much have you budgeted for EU AI Act implementation efforts?
Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

Some have implemented GDPR-related controls, technical guardrails, and model fine-tuning for external models. But just 16 percent of respondents are conducting red-teaming efforts, while some said they are rolling back relationships with suppliers while rules and obligations for general-purpose AI become applicable throughout 2025.

Data governance. Only 18 percent of respondents said their organizations have mature technical risk management processes for AI systems in place. In addition, few have robust models or security and accuracy techniques. However, about 75 percent of respondents indicated they had advanced cyber controls and data protection measures in place.

The act introduces requirements for data management. These cover choices in designing systems, formulating assumptions, collecting and preparing data, examining bias, ensuring representative data use, and including the appropriate statistical properties. More than half of survey respondents said they have not yet addressed these requirements. Less than 20 percent have addressed bias.

What models do with the data is another area of concern. Many respondents cited difficulty in defining standards for testing the outputs of gen AI models. For self-developed models, respondents said they commonly use continuous code integration and deployment, model versioning, and documentation to ensure quality.

Thirty-eight percent of respondents use "human in the loop" processes, while 30 percent use technically responsible AI tooling. Model performance monitoring, logging, and user feedback, together with incident detection and management, are the most common measures used to ensure quality after deployment.

Talent. Getting the right people to run and manage AI is proving difficult, too. The talent shortage is especially prominent for technical staff but also exists for legal personnel. This is a major concern not only for businesses but also for regulatory authorities that have concerns about competent monitoring and enforcement of the AI Act. Only a quarter of respondents upskill employees, which takes time and investment.

Other. Perhaps surprisingly, respondents did not cite cost, financial resources, or ethical concerns as top reasons for the slow progress on implementation (Exhibit 6).

Given the complexity of the EU AI Act and the effort needed to comply, it would be prudent for organizations to accelerate their planning now.

Exhibit 6

Key challenges of implementing the EU AI Act relate to unclear obligations, complexity, and talent gaps.

Key challenges facing organizations in implementing the EU AI Act,[1] % of respondents

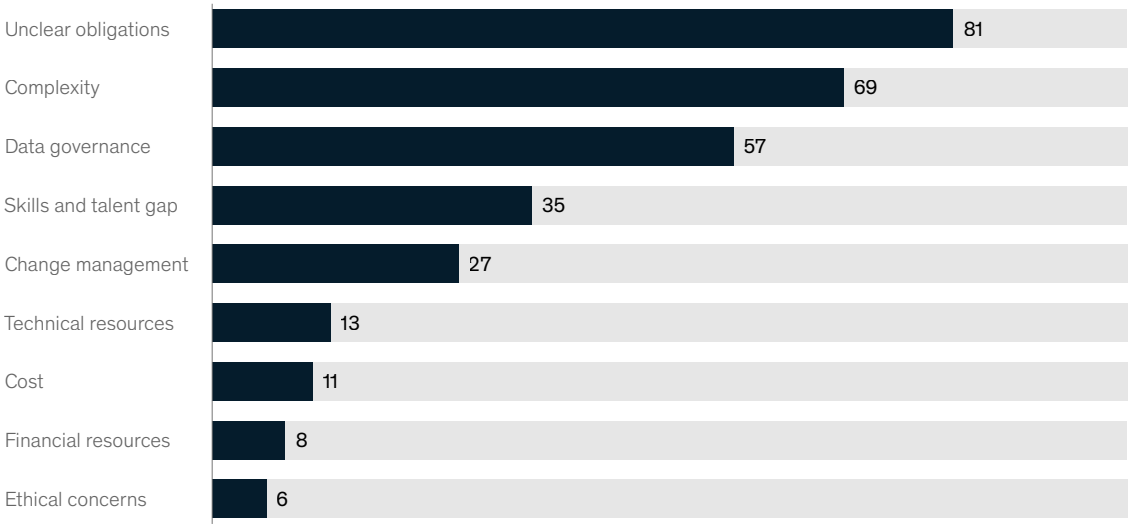

Source: McKinsey EU AI Act Survey, spring 2024 (n = 180 organizations in Europe)

McKinsey & Company

The time to act

Given the complexity of the EU AI Act and the effort needed to comply, it would be prudent for organizations to accelerate their planning now. While the act outlines implementation stages and staggered compliance deadlines, those with experience implementing GDPR understand that waiting can create chaos as those deadlines approach.

Managing the scope of an organization's AI efforts is important. Organizations that align development to governance practices manage to limit the number of models they use, generally to fewer than 20. A clear governance structure can also limit teams' frustrations in fielding ad hoc requests and trying to get support.

Organizations should embrace a "define your world" approach, which prioritizes transparency in model use, stakeholders, risks, and regulations. The EU AI Act has set out requirements mainly for high-risk models, so a risk categorization of the model landscape will help structure the work going forward and control the level of effort.

Defining a target state for governance and compliance efforts can help organizations build road maps that thoroughly consider strategy, risk appetite, organizational structure, technology, policy, and tooling. And organizations can continue to get better through a process of ongoing improvement, using existing best practices and frameworks as a guide. Ensuring cross-functional collaboration and input on ethical and risk considerations is paramount, so if current risk

functions are not equipped, separate action on top of existing governance may be required.

To achieve compliance, organizations will need the necessary talent, resources, and relevant KPIs to measure progress. AI is evolving quickly, so it is essential to stay on top of changes. The EU AI Act represents a significant step toward regulating AI systems and ensuring responsible AI governance and could serve as a blueprint for other jurisdictions globally.

But before that happens, the act's regulators will need to further clarify their expectations and work with the industry to find pragmatic implementation solutions in an environment of limited resources. Responsible and trustworthy AI is a prerequisite to defining a new digital future. By embracing responsible AI governance, companies can spur innovation with the trust of consumers, competitors, shareholders, and society behind them.

This article originally appeared in the August/September edition of The RMA Journal.

Henning Soller is a partner in McKinsey's Frankfurt office; **Anselm Ohme** is a consultant in the Berlin office, where **Chris Schmitz** is a data science fellow; **Malin Strandell-Jansson** is an alumna of the Stockholm office; **Timothy Chapman** is an analyst in the Wroclaw office; and **Zoe Zwiebelmann** is a consultant in the Hamburg office.

The authors wish to thank Andreas Kremer, Angela Luget, Angie Selzer, Artem Avdeed, and Silvia Tilea for their contributions to this article.

Designed by McKinsey Global Publishing