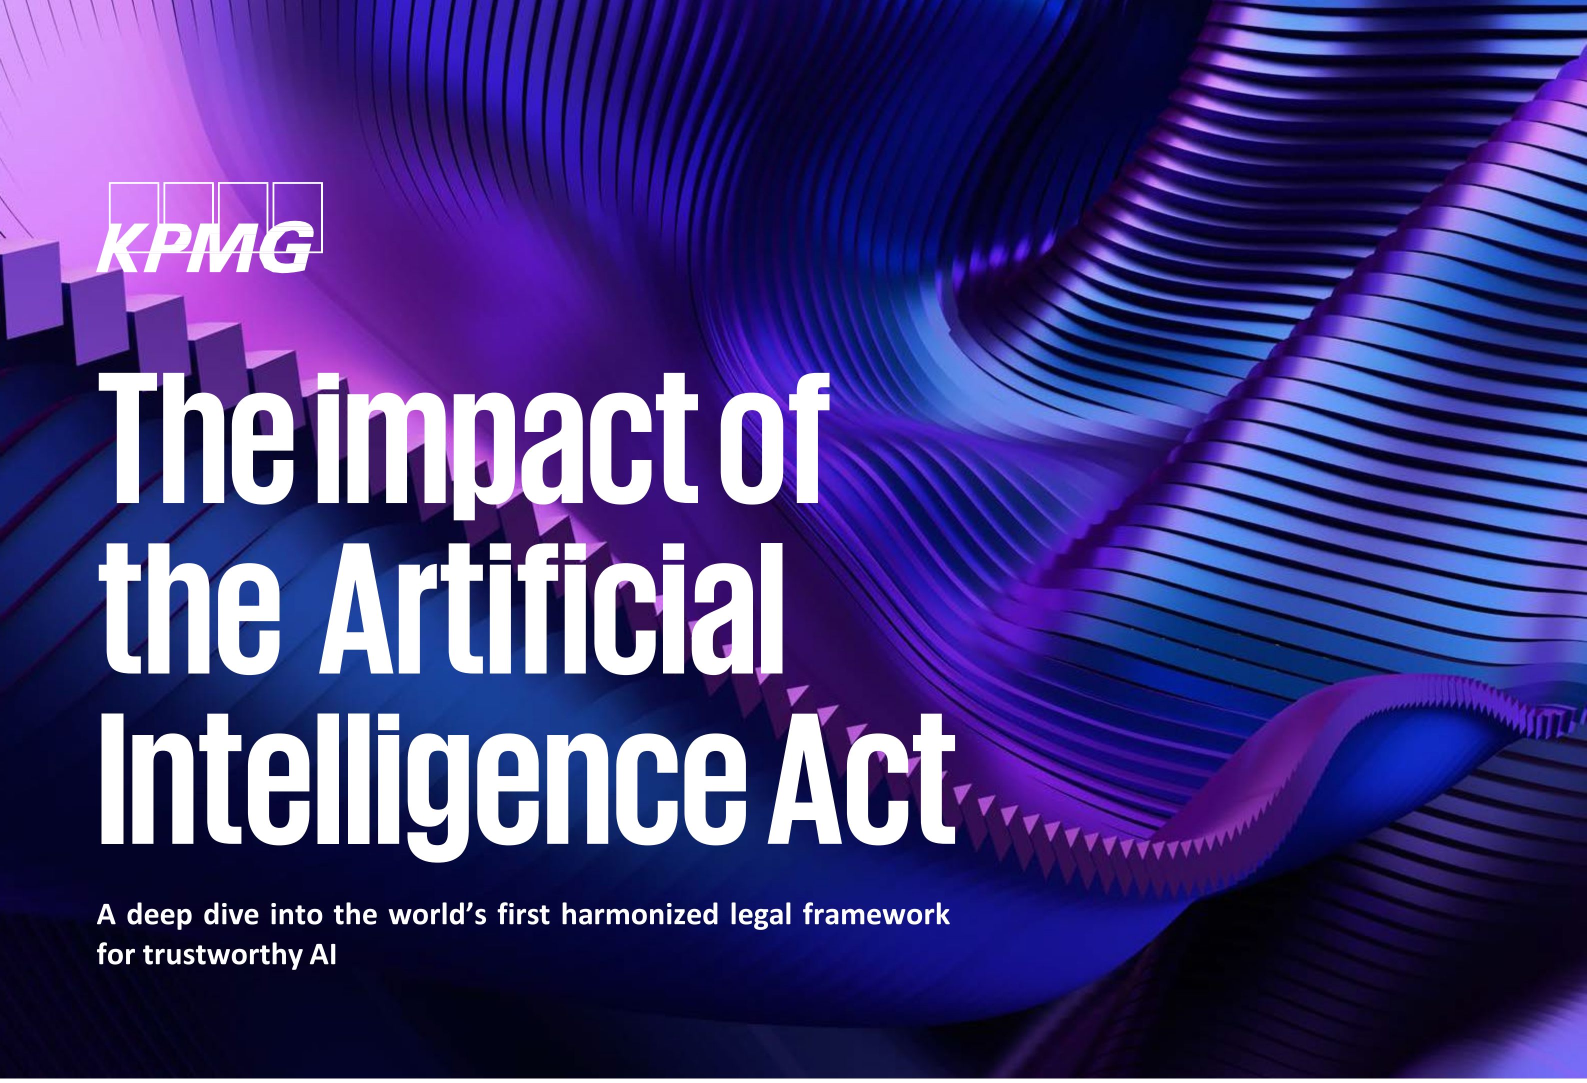

The impact of the Artificial Intelligence Act

A deep dive into the world's first harmonized legal framework for trustworthy AI

Artificial Intelligence (AI) encompasses an evolving range of technologies that are increasingly used to tackle societal and industry problems. While AI offers obvious values in a multitude of sectors, the technologies that drive these advantages also entail certain risks and potential negative consequences. The EU reached agreement on a comprehensive AI Act, which should ensure protection of public interest and fundamental rights of citizens in the rapidly evolving AI landscape.

On Friday December 8th, the European Union reached a provisional agreement on the AI Act, after the longest negotiations in EU history. The AI Act will be the first of its kind and will most likely become a global standard for AI regulation. In this paper, we take a deep dive into the AI Act, the timelines, its structure, and the obligations it introduces. While the final text is not yet publicly available, we take you along what is known and to be expected.

With the introduction of the AI Act, the European Union aims to strike a balance between fostering AI development and uptake on the one hand, and maintaining public interest and fundamental rights protection on the other.

Most obligations are expected to come into effect as early as the first half of 2026. Prohibitions in the AI Act are anticipated to take effect by the end of 2024, and obligations regarding general-purpose AI are expected as early as 2025. The proposed regulation will have a far-reaching impact on all organizations leveraging this technology. The consequences of non-compliance could range from restriction of market access to the product, to administrative fines of up to 7% of annual turnover.

Due to the broad definition of AI that is expected to be used in the AI Act, most organizations are developing or using systems that will qualify as AI under this regulation. Given the short implementation period, it is imperative that organizations start gaining a thorough understanding of the AI systems they develop and deploy, in order to gain insight into the impact of this new regulation on their business and operations.

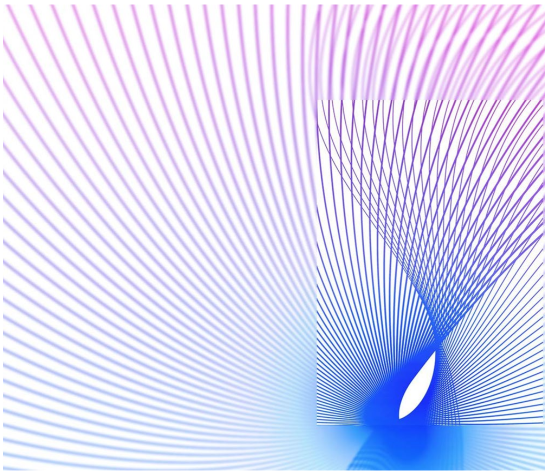

> **The AI Act requires responsible management and transparency**

The Artificial Intelligence Act on a page

A summary

The AI Act aims to balance AI uptake and protection of citizens

Artificial Intelligence (AI) systems are being deployed to address various societal and industrial challenges; this, however, also brings about certain risks and potential negative consequences. In order to regulate and mitigate these concerns and to balance the benefits of AI with regulatory measures, the European Union introduces the AI Act. A political agreement on the AI Act has been reached and the final text is expected to be available in the first half of 2024.

The objective of the AI Act is to ensure that AI systems are safe, respect fundamental rights, provide legal certainty for investment and innovation, enforce safety requirements and fundamental rights effectively, and prevent market fragmentation by creating a single market for trustworthy AI application.

Most organizations are building or using systems that need to comply with the AI Act as early as the first half of 2026

The definition of Artificial Intelligence under the AI Act is expected to be broad. This definition is likely to cover a wide range of technologies and systems. Due to this broad definition of AI and due to the nature of the systems that fall in the high-risk category, almost all organizations develop or deploy high-risk AI systems. This broad scope and the extent of obligations, mean that the AI Act will have far-reaching impact on many aspects of organizations' operations and management.

Several of the obligations are expected to apply as early as the first half of 2026 and will directly impact all organizations who build or deploy AI systems in EU markets.

> "It's up to organizations' leadership to drive Responsible AI in line with the AI Act, company brand, values and risk tolerance"

Developers of high-risk systems will have to comply with the most stringent obligations

To achieve its objectives, the AI Act applies a risk-based regulatory approach, dividing AI systems into four categories: unacceptable risk, high risk, limited risk and minimal risk.

For high-risk AI systems, the AI Act imposes the most stringent obligations. These obligations will primarily affect developers of AI systems, called 'provider' in the AI Act. Developers will need to ensure that their AI systems comply with strict standards concerning risk management, data quality, transparency, human oversight, and robustness, in order to minimize risks. The term 'provider' includes organizations that develop AI systems for use within their own organization.

Users of AI systems face obligations with regard to responsible use

Users of high-risk AI systems also face new regulatory obligations. Users are responsible for operating these AI systems within the legal boundaries set forth by the AI Act and in line with instructions for use set out by the developer. This includes obligations on data handling, human oversight and monitoring. An organization can be both a user and a developer.

Understanding the impact of the AI Act on your organization starts with insight into your portfolio of AI systems

Organizations should take the time to create an overview of AI systems they develop and/or use and map these to the risk levels defined in the AI Act. If any of their AI systems fall into the limited, high or unacceptable risk category, they will need to assess the impact the AI Act has on their organization. It is imperative to understand this impact as soon as possible. Prohibitions are expected to apply six months after final adoption, which means they will likely apply from the end of 2024. Obligations on high risk AI systems will likely apply as early as the first half of 2026.

The AI Act takes a risk-based approach

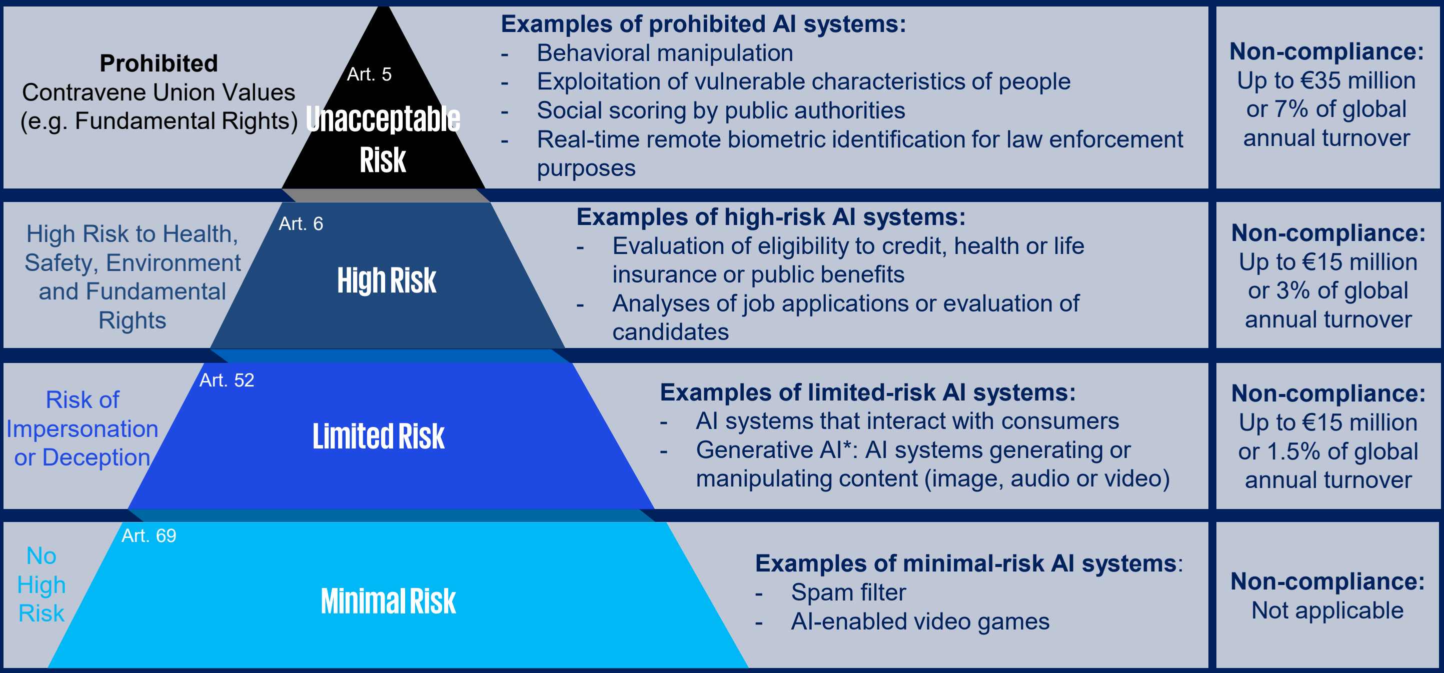

Prohibited Contravene Union Values (e.g. Fundamental Rights)	Art. 5 Unacceptable Risk	**Examples of prohibited AI systems:** - Behavioral manipulation - Exploitation of vulnerable characteristics of people - Social scoring by public authorities - Real-time remote biometric identification for law enforcement purposes	**Non-compliance:** Up to €35 million or 7% of global annual turnover
High Risk to Health, Safety, Environment and Fundamental Rights	Art. 6 High Risk	**Examples of high-risk AI systems:** - Evaluation of eligibility to credit, health or life insurance or public benefits - Analyses of job applications or evaluation of candidates	**Non-compliance:** Up to €15 million or 3% of global annual turnover
Risk of Impersonation or Deception	Art. 52 Limited Risk	**Examples of limited-risk AI systems:** - AI systems that interact with consumers - Generative AI*: AI systems generating or manipulating content (image, audio or video)	**Non-compliance:** Up to €15 million or 1.5% of global annual turnover
No High Risk	Art. 69 Minimal Risk	**Examples of minimal-risk AI systems**: - Spam filter - AI-enabled video games	**Non-compliance:** Not applicable

* Further specific obligations to generative AI and foundation models will apply outside of this risk-based approach.

A deep dive into the Artificial Intelligence Act

Legislative procedure of the AI Act

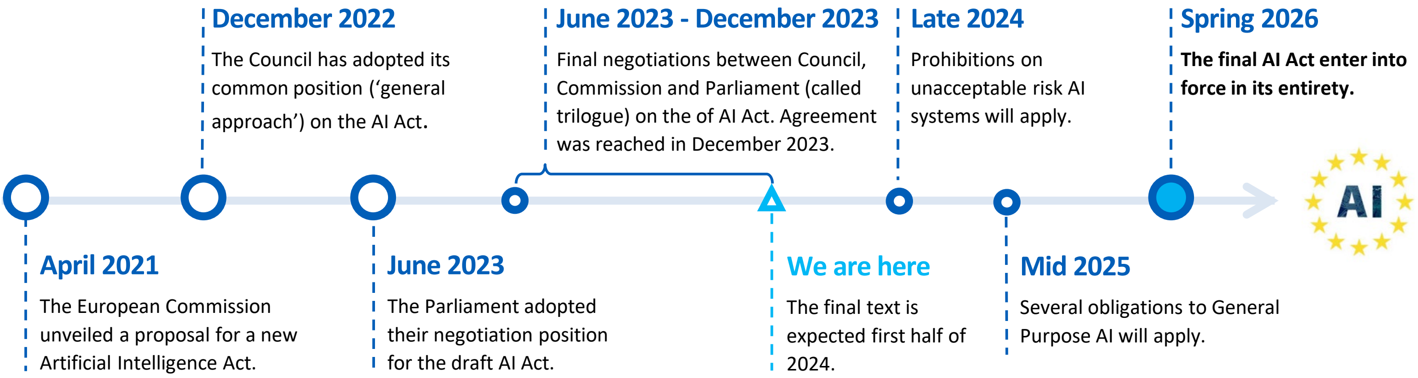

April 2021
The European Commission unveiled a proposal for a new Artificial Intelligence Act.

December 2022
The Council has adopted its common position ('general approach') on the AI Act.

June 2023
The Parliament adopted their negotiation position for the draft AI Act.

June 2023 - December 2023
Final negotiations between Council, Commission and Parliament (called trilogue) on the of AI Act. Agreement was reached in December 2023.

We are here
The final text is expected first half of 2024.

Late 2024
Prohibitions on unacceptable risk AI systems will apply.

Mid 2025
Several obligations to General Purpose AI will apply.

Spring 2026
The final AI Act enter into force in its entirety.

Objectives of the AI Act

The EC aims to strike a balance between fostering AI development and uptake, and maintaining public interest and fundamental rights protection. This is reflected in the objectives of the proposal:

- ensuring that AI systems placed and used on the EU market are safe and respect fundamental rights and EU values;
- ensuring legal certainty to facilitate investment and innovation in AI systems;
- enhancing governance and effective enforcement of fundamental rights and safety requirements that apply to AI systems;
- facilitating the development of a single market for lawful, safe and trustworthy AI applications and preventing market fragmentation.

To achieve these objectives, the AI Act applies a risk-based approach. This allows for establishing certain minimum necessary requirements to address the risks and problems linked to AI systems, without unduly constraining or hindering technological development or disproportionately increasing costs relating to placing AI systems on the market.

Definition of AI and the scope of the Act

The AI Act definition of AI is in line with the definition of AI by the OECD, with some (minor) changes. The definition text is not yet publicly available, but the OECD definition is as follows:

"An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."

The definition of AI in the Act ultimately defines its scope. This definition extends beyond impressive applications of deep learning and generative AI and includes far simpler technologies and systems. As a result the scope of the act is much broader than was expected based on early proposals, extending significantly beyond a conventional understanding of AI.

There will be several exceptions which will be out of scope for the AI Act. There are exemptions for AI systems used for military or defense purposes, as well as limited exemptions for free and open source systems.

The AI Act in a nutshell

Stimulate the good

- Stimulate innovation through regulatory sandboxes
- Stimulate harmonization of standards, codes of conduct and certification
- Offer greater transparency regarding AI systems
- Create level playing field for actors involved
- Safeguard fundamental rights and provide legal certainty for EU citizens

Fix the bad

- Impose stricter requirements for high-risk AI systems (obligatory risk management, data governance, technical documentation, etc.)
- Carry out conformity assessments and post-market monitoring for high-risk AI systems
- Avoid fundamental rights violations
- Establish effective oversight and enforcement mechanisms

Control the ugly

- Prohibit unacceptable-risk AI systems
- Prevent use of subliminal techniques that distort a person's behavior in such a way that it causes harm to that person or another person
- Prohibit exploitation of vulnerabilities of a specific group of persons, e.g. exploiting age or disability.

The AI Act will apply to all AI systems built or deployed in EU markets.

The risk-based approach explained

The AI Act makes a distinction between AI systems based on the risks that are associated with them on the basis of how the systems are used. 'High risk' is viewed through the lens of the impact on people and is independent from the systems' complexity. The obligations for the actors involved will depend on the category of AI systems at hand. Although an agreement on the context has been reached, the final text is not yet available. The following sections summarize the obligations of the AI Act based on publicly available information.

Unacceptable-risk AI systems

A. What kinds of AI systems are covered?

AI systems that enable manipulation, exploitation and social control practices are seen as an unacceptable risk.

This category would prohibit systems for the use of:

- Manipulation in such a way that it causes (or is likely to cause) harm to that person or another person;
- Exploiting vulnerabilities of a specific group of persons
- Social scoring leading to detrimental or unfavorable treatment in social contexts;
- Indiscriminate scraping of facial images;
- Emotion recognition software in the workplace and education (with some exceptions);
- Use of AI systems that categorize persons based on sensitive traits (e.g. race, political opinions, or religious beliefs);
- Predictive policing on individuals (risk scoring for committing future crimes based on personal traits);
- Remote biometric identification of natural persons (partial ban with some exceptions in law enforcement).

This list is more extensive than it was in the original proposal. The category of prohibited AI systems was subject to heavy negotiations in the last trilogue and not all outcomes of this discussion are available to the public. Details in the final text will provide a deeper understanding of both the exact prohibited categories and exceptions.

B. What are the obligations related to this category?

Since the AI systems in this category are deemed an unacceptable risk, **they are prohibited**.

High-risk AI systems

A. What kinds of AI systems are covered?

The AI Act establishes a list of categories of AI systems that are considered high risk. It includes the following uses of AI systems:

- Biometric identification systems (those not already prohibited under "unacceptable risk");
- Critical infrastructure (e.g. supply of utilities).
- Educational and vocational training (e.g. automated scoring of – or exclusion from – exams);
- Employment, workers management, and access to self-employment (e.g. automated recruitment and application triage);
- Access and enjoyment of essential private and public services (e.g. benefits systems, credit, insurance);
- Law enforcement systems that may interfere with fundamental rights (e.g. automated risk scoring with regard to potential offenders, deepfake detection software, evidence reliability scoring);
- Migration, asylum and border control management (e.g. verification of authenticity of travel documents; visa and asylum application examination);
- Administration of justice and democratic processes (e.g. legal interpretation tools to assist judicial authorities).

Most organizations use these high-risk AI systems, such as AI systems for recruitment purposes.

Additionally, the high-risk category consists of AI systems that satisfy both of the following requirements:

- they are (intended to be) used as a safety component of a product, or are themselves a product, covered by an exhaustive list of EU harmonization legislation; *and*
- a third-party conformity assessment is required pursuant to the aforementioned EU harmonization legislation.

In practice, the products covered are, among other things: machinery, toys, medical devices, multiple kinds of vehicles, marine equipment, and lifts.

B. What are the obligations related to this category?

Since the AI systems in this category are considered to be high risk, they are subject to the most stringent regulatory requirements:

- Adequate risk management system to identify, evaluate and mitigate risks during the lifecycle of the AI system;
- Appropriate data governance and management practices (training, validation and testing) need to be implemented to ensure the quality of the datasets;
- Technical documentation to be drawn up to demonstrate compliance with obligations and allow for compliance assessments;
- Logging of events to ensure traceability of the system's functioning;
- Registration in EU Database for high-risk AI systems;
- Transparency obligations to enable correct interpretation and usage of the AI system;
- Implementation of appropriate human oversight measures;
- Appropriate levels of accuracy, robustness and security to be achieved.

High-risk AI systems will be subject to prior conformity assessment procedures to determine whether they comply with the aforementioned requirements. As a final step before being placed on the market, a declaration of conformity must be signed and the AI system must be provided with the CE marking, although the specifics of such standards are not yet clear. Exceptions exist for law enforcement when transparency obligations may conflict with public safety.

Once the AI system is put on the market, post-market monitoring obligations will apply. This includes reporting serious incidents/malfunctioning of high-risk AI systems to the relevant market surveillance authorities.

General purpose AI, foundation models and generative AI

A. What kinds of AI systems are covered?

'General purpose' and 'foundation model' AI were not defined in the original proposal, but have been included in the final Act after heavy negotiations.

- General purpose AI systems are AI systems that are intended to perform generally applicable functions (such as image/speech recognition, audio/video generation, pattern detection, etc.) and that can be used in and adapted to a wide range of applications. Well-known examples include 'generative AI' applications such as ChatGPT and Dall-E.
- Foundation models are AI systems that are trained on broad data at scale, designed for generality of output, and can be adapted to a wide range of distinctive tasks. A well-known example is GPT-4, which is the foundation model under the latest ChatGPT.

B. What are the obligations related to this category?

General purpose AI systems will have to comply with transparency requirements. These include technical documentation, complying with EU copyright law and providing information on the training data.

For the most powerful foundation models more stringent obligations will apply. Powerful is defined by a set boundary in computing power. Providers of these models will have to conduct model evaluations, assess and mitigate systemic risks, conduct adversarial testing, report to the Commission on serious incidents, ensure cybersecurity and report on their energy efficiency.

Limited-risk AI systems

A. What kinds of AI systems are covered?

Some AI systems that are intended to interact with natural persons or generate content would not necessarily qualify as high-risk AI systems, but may nevertheless entail risks of impersonation or deception. This includes the outputs of most generative AI systems.

In practice, the following AI systems are to be identified in this category:

- Chatbots, such as ChatGPT-based systems
- Emotion recognition systems
- Biometric categorization systems
- Systems generating 'deepfake' content

B. What are the obligations related to this category?

AI systems in this category should comply with certain transparency obligations. Unlike the obligations for high-risk systems that impact development and risk management in a broad sense, the obligations for limited-risk systems focus on outputs and users:

- Natural persons must be informed that they are interacting with an AI system;
- Natural persons exposed to a (non-prohibited) emotion recognition or biometric categorization system must be informed about the operation of the system;
- 'Deepfake' content must be disclosed as being artificially generated or manipulated.

Minimal-risk AI systems

A. What kinds of AI systems are covered?

This category of AI systems is not defined by the AI Act. These AI systems are the systems that are not included in the other categories described above. This category includes applications such as AI-enabled video games or spam filters.

B. What are the obligations related to this category?

In the AI Act, this category of AI systems will not be subject to stringent obligations, with the exception of adhering to general product safety standards. Nevertheless, the promotion of establishing codes of conduct is strongly encouraged, where the assumptions is that doing so holds the potential to foster a wider adoption of reliable Artificial Intelligence within the EU.

A deep dive into the requirements for high-risk AI systems

The possible use cases of high-risk systems should be viewed broadly. Departments potentially deploying a high-risk AI system include HR, Finance, Customer Services, IT and Legal. High-risk AI systems will be subject to the following predefined requirements:

Risk Management

- A risk management system must be implemented, documented and maintained;
- The establishment of a risk management system is a continuous iterative process;
- Obligations include identifying, evaluating, and managing potential risks through suitable measures that are regularly updated and tested.

Data and data governance

- Requirements for data governance and management practices in the development of high-risk AI systems making use of training techniques involving data, must be met.

Technical documentation

- The technical documentation must be drawn up before the AI system is placed on the market or put into service, and be kept up to date;
- The technical documentation must demonstrate compliance with requirements, be kept up to date, and include minimum elements as listed in Annex IV of the AI Act.

Record-keeping

- High-risk AI must log events to trace and monitor high-risk situations, conforming to standards;
- Minimum logging must include usage, data, and personnel identification.

Transparency and provision of information to users

- Operations of AI systems must be sufficiently transparent to enable users to interpret the system's output and use it appropriately;
- Therefore, AI systems must be accompanied by instructions in an appropriate digital format.

Human oversight

- Measures for effective human oversight of high-risk AI systems must be in place, including appropriate tools, monitoring, interpretation, and intervention.

Accuracy, robustness and cyber security

- Requirements related to appropriate accuracy, robustness, cyber security levels, resilience to errors, biases, unauthorized access and more, must be met.

AI Act and GDPR

The AI Act and the GDPR should be seen as complementary frameworks. Each comes with its own set of rules and obligations, and, in practice, both will often need to be applied simultaneously. Many AI systems will process personal data in the context of their operations, therefore requiring compliance with both the future AI Act and the GDPR. Even though compliance of AI systems with some GDPR principles (e.g. purpose limitation and data minimization) might prove challenging in practice, already having in place the necessary data protection controls and policies will prove to be an advantage to those organizations who also engage in the development or use of AI systems.

Some parallels that can be identified between the AI Act and GDPR are indicated in the table below:

AI Act	GDPR	Comment
Risk-based approach: obligatory risk management system for high-risk AI and Fundamental Rights Impact Assessments	Risk-based approach: data protection impact assessments (DPIAs) for high-risk processing activities	Transparency information of high-risk AI systems to be used in DPIAs under GDPR
Transparency obligations for most AI systems	Transparency obligations regarding personal data	Offering transparency with regard to AI systems may help to achieve transparency as required in the GDPR
Robustness and cybersecurity obligations for high-risk AI, taking into account relevant risks	Appropriate technical security measures to protect personal data, taking into account relevant risks	Security measures relating to AI systems may benefit the protection of personal data and vice versa
Reporting serious incidents/malfunctioning of high-risk AI systems to market surveillance authorities	Reporting personal data breaches to data protection authorities	Reporting obligations relating to serious incidents/malfunctioning of AI systems may partially overlap with GDPR reporting obligations when personal data is involved
Providers/users must be able to demonstrate compliance of high-risk AI systems with relevant requirements	Controllers must be able to demonstrate compliance with GDPR principles (accountability)	Having procedures in place to document and register the actions taken to demonstrate compliance is beneficial in the context of both frameworks
Administrative penalties for non-compliance with the Act	Administrative penalties for non-compliance with the Act	Infringements give way to far-reaching penalties under both frameworks

What does the AI Act mean for your organization?

Almost all organizations develop or deploy high-risk AI systems. AI is expected to be broadly defined under the AI Act and the definition is likely to cover a wide range of technologies. Additionally, the nature of systems that fall in the high-risk category covers a wide range of applications that many organization regularly apply, such as AI systems for recruitment purposes. This broad scope and the extent of obligations means that the AI Act will have a far-reaching impact on many aspects of organizations' operations and management. As most organizations are developing AI at pace, strong and adaptable guardrails are needed to keep up with development and the AI Act's requirements.

The consequences of non-compliance could range from restriction of market access to the product, to administrative fines of up to 7% of annual turnover; a fine which even exceeds that of the GDPR.

The AI Act could enter into force as early as the first half of 2024. Several obligations to general purpose AI systems may already apply in 2025. Other obligations are expected to come into effect as early as the first half of 2026. This means organizations have a window of 2 years to organize their operations and processes to become compliant for their high-risk systems.

What's next?

The final text of the AI Act is expected to be published first half of 2024. Because of the short implementation periods you can proactively take steps to ensure your organization is well prepared.

A first step is to ensure the right people in your organization start working on these upcoming regulatory requirements as soon as possible. Early engagement ensures you understand the requirements and their impact. The AI Act identifies a range of roles that includes Legal, Privacy, Data Science, Risk Management and Procurement professionals. Therefore a multidisciplinary taskforce responsible for compliance with the AI Act should include this range of expertise.

Secondly, it is crucial to gain a comprehensive understanding of AI systems developed or used in your organization and to map these to the risk levels defined in the AI Act. If it turns out that any of your AI systems fall into the limited risk, high risk, or unacceptable risk category, the AI Act requires impactful changes to processes and operations before 2026. It is imperative to understand this impact as soon as possible, in order to manage the required changes and ensure timely compliance with the new legal framework when it comes into effect.

KPMG's interdisciplinary AI expertise

KPMG has been building knowledge and experience in AI technology and the responsible use of AI for years. With an interdisciplinary approach, we possess the required expertise to tackle a diverse array of potential use cases across various industries. Not only has KPMG developed a responsible AI methodology, but also specific tools, practices and ways of working to attain practical solutions, enabling clients to be compliant to AI-related laws and regulations while preserving flexibility and innovation with this technology.

The extensive obligations on high-risk AI systems are impactful but not entirely new. KPMG's multidisciplinary teams, consisting of Data Science and AI specialists, (IT) Risk Management, Digital Transformation and Tech Law professionals, have built extensive experience supporting our clients with the establishment of risk management systems for AI, the implementation of data governance and evaluations of human oversight, and accuracy and robustness in deployed systems, among other topics. In our work, we cover the full spectrum of AI applications for various clients. From credit risk models, to content moderation, to generative AI; from government, to insurers, to global tech organizations.

Our methodology builds on the most recent standards (e.g. NIST and ISO standards), industry better practices, and a deep understanding of AI-related laws and regulations, whether those are general or sector-specific, Europe-wide or national.

We offer a comprehensive approach that enables your organization to effectively manage these upcoming regulatory changes. Through our services, we assist you in embarking on your transformation and compliance journey aligned with and customized to your business requirements.

Contact our specialists

Ylja Remmits
Sr. Consultant Responsible AI
+31 (0) 6 30844868
Remmits.Ylja@kpmg.nl

Frank van Praat
Director Responsible AI
+31 (0) 6 51206152
vanPraat.Frank@kpmg.nl

Peter Kits
Partner Digital Law
+31 (0) 6 13001055
Kits.Peter@kpmg.nl

Sander van der Meijs
Director AI strategy & transformation
+31 (0) 6 52078891
vanderMeijs.Sander@kpmg.nl

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.