Privacy in AI: Enhancing Innovation

The better the question. The better the answer.
The better the world works.

Shape the future
with confidence

In brief

- The importance of privacy within AI can be safeguarded through the interplay between the GDPR and the AI Act.
- The AI Act and the GDPR differ significantly in their purpose and scope, which may result in potential conflicts between key principles of both regulations.
- Discover a proactive and practical approach to ensure compliance and protect privacy

Introduction

Artificial Intelligence (AI) is rapidly advancing, yet the safeguards to protect fundamental human rights, such as personal privacy, are still insufficient. The impact of AI on privacy cannot be overstated, especially since we do not fully understand its implications. Without proper regulations, AI could be misused, leading to significant breaches of privacy and human rights. Just as brakes are essential in cars, enabling us to drive faster and more safely by providing control, robust legislation is crucial for AI. Effective regulations and authorities ensure that AI development and deployment protect these rights. By implementing safety measures, we can harness AI's potential to foster a society where technology enhances our lives while upholding our values.

The best way to discuss preserving privacy in AI and protecting fundamental rights is by looking into regulations governing such aspects. Since 2018, personal data privacy has been safeguarded under the European General Data Protection Regulation (GDPR). According to this legislation, personal data refers to "any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person." Building on this foundation, the EU Commission introduced the AI Act (AIA), set to be fully effective in 2026. The AIA aims to protect human rights, such as privacy, ensuring that AI technologies are developed and used in ways that respect them. It defines AI as "a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and 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." AI models rely on a vast amount of data, including personal data, to function effectively. As a result, certain unavoidable tensions between the GDPR and AIA are building up which need to be addressed.

The interplay between the GDPR and the AI Act

The AIA adopts a horizontal approach to influence AI standards beyond its borders. Drawing from product safety regulations, it employs a risk-based approach that assesses risks on a broader scale, including impacts on social groups and individuals not directly connected to data processing. However, this approach raises legal questions about its alignment with the GDPR, given their differing conceptions. Non-compliance with AIA requirements will result in penalties similar to those under the GDPR, including cease-and-desist orders and market withdrawals.

In contrast, the GDPR adopts a principled, individual-centric approach based on the fundamental right to personal data protection outlined in Article 8 of the European Charter of Fundamental Rights. It establishes principles, grants rights to individuals, and imposes obligations on organizations collecting data, with varying obligations based on the risk associated with the data processing. Data controllers and data processors that fail to comply with these requirements will face penalties imposed by the supervisory authority.

When comparing both legal instruments, several key challenges can be identified:

1. GDPR principles (Art. 5) and AI systems development and use

AI's reliance on vast amounts of data, including personal data, raises concerns about compliance with GDPR principles like data minimization and purpose limitation. Ensuring compliance is challenging when AI providers use data collected for different purposes or from third parties with alternative objectives. The AIA's requirements for high-risk AI systems to ensure training data relevance, representativeness, error-free, and completeness may conflict with Article 5 of the GDPR, which contains the main principles of the legislation.

The GDPR's first principle is lawfulness, fairness, and transparency. Any data that is processed must be done in a transparent, lawful, and fair manner. When applying this principle to the AIA and AI systems, multiple questions remain unanswered. For instance, an AI system based on personal data must be developed with a clearly defined "purpose" or objective. This helps to frame and limit the data used for training, ensuring that only relevant data is processed. The purpose must be established as soon as the project is defined. Such a purpose must be explicit, known, understandable, and legitimate, aligning with the organization's tasks. While the need for a well-defined purpose may conflict with the potential for AI systems to generate general-purpose solutions, this requirement can be adapted to the context of AI without being disregarded.

There are two types of situations when developing an AI system: one where the developer knows the use of the AI system, and another where the developer cannot clearly define the use of the AI system. In the former, the purpose is established in the development phase, explicit and legitimate regarding the identified operational use. In the latter, it gets more complex. When developing an AI system that can be used in different contexts and has multiple applications (e.g., ChatGPT), the purpose of the AI system is defined more generally. A good practice is to be as precise as possible and refer to the type of system being developed (e.g., GenAI for image creation) and its foreseeable functionalities and capabilities.

By defining the purpose, we can establish the lawfulness of processing. We need to determine the most suitable legal basis for the given situation. Consent is often considered the most appropriate legal basis, as individuals have the freedom to accept or refuse to give away their data without facing negative consequences (such as losing access to the service). As stated in the legislation, consent must be freely given, specific, informed, and unambiguous. However, obtaining consent is often impractical or even impossible for dataset creation or model training for multiple reasons. Legitimate interest is another possibility. Using legitimate interest can only be considered if the following three requirements are met: (1) the interest pursued is legitimate, legal, precisely, and genuinely defined; (2) it must be possible to establish that the personal data is necessary for the training of the system, as relying on non-personal or anonymized data is not a viable option; (3) the use of such personal data must not lead to "interference" with the privacy of individuals. Other legal bases, such as contractual or legal obligations, may be used more exceptionally. In these cases, you must demonstrate how

your processing is necessary to fulfill the contract, pre-contractual measures, or a sufficiently precise legal obligation.

Transparency is crucial regarding AI systems. These systems tend to be black boxes, making the understanding of the outcome complicated or even impossible, thus making it difficult to be transparent towards data subjects.

Another important principle that must be addressed is purpose limitation. AI systems training requires the use of as much data as possible to increase the accuracy of the system and limit the level of bias and the risk of discrimination, unfair treatment, etc. However, the use of data is limited by the 'purpose limitation' principle, which requires companies to collect data for specified, explicit, and legitimate purposes and not further process it in a manner that is incompatible with those purposes.

Data minimization is closely related to this principle. AI system training requires the use of as much data as possible, whereas the 'data minimization' principle requires companies to limit the use of data to what is relevant and necessary for the purposes for which they're processed. This is where issues may arise. AI systems may not always respect the principle of data minimization, as they rely on a large quantity of data and do not consider the relevance or necessity of it. This principle is in direct conflict with AI system development and use.

In addition to data minimization, the accuracy of the data and storage limitation must also be considered. Data needs to be correct, and inaccurate data needs to be erased or rectified without delay. However, AI systems tend to produce results for which it's not possible to assure accuracy, bringing forth another potential limitation and breach. Concerning storage limitation, the data used to train an AI system ideally needs to be kept for reasons such as re-training the AI system with additional data or supporting transparency obligations.

Lastly, the principles of integrity and confidentiality, in addition to accountability, must be met. System training often requires the use of personal data in non-production environments, where security is rarely at the same level as in production environments.

Due to the tensions present with the above GDPR principles, it becomes more difficult and complicated for a data controller to be accountable for the processing of personal data and ensuring compliance with the GDPR. As a result, companies will need to take a position on the above issues without being assured that this position will be supported by the supervisory authority and courts. Therefore, companies should document and justify their decisions and regularly check their decisions based on future guidelines and decisions of the supervisory authority and courts.

2. GDPR's view on automated decision making in Article 22

Under Article 22 GDPR, individuals have the right not to be subject to decisions made solely by automated systems. Individuals can invoke this right if these decisions have significant legal or similar consequences, such as affecting one's rights, status, or access to services. When automated decisions are made, they must meet specific conditions. For instance, some conditions can be that it must be based on the individual's explicit consent, that it is a contractual necessity, or that it is authorized by law.

The challenge here arises when defining what constitutes "human oversight" and when a decision is classified as fully automated. Article 22 GDPR does not prohibit all automated decision-making but requires human intervention in certain cases to ensure fairness and accountability. For example, the inclusion of a human at a critical stage of the decision-making process may shift the decision from being classified as fully automated. The moment when this intervention occurs is important. Mere oversight is not sufficient; human involvement must be meaningful and impactful enough to influence the outcome. The AIA further emphasizes the importance of human intervention, particularly in high-risk AI systems, by mandating that such systems be designed with tools to facilitate human monitoring and control, known as the "human-in-the-loop" approach. This supervision can involve regular checks or adjustments by humans to ensure the system operates within safe, ethical, and lawful parameters.

3. Data subject rights vs. AI model integrity

The GDPR grants individuals' rights such as access to their data, correction of inaccuracies, and deletion of data. AI models rely on large datasets, and data deletion requests could affect the integrity and performance of AI systems. It will be challenging to reconcile individuals' rights to control their data with the operational requirements of AI systems.

4. Inference, resulting in the expansion of the GDPR's scope

AI's advanced inference and correlation capabilities may identify individuals without using their data in training sets, expanding the GDPR's scope. AI systems trained on personal data can generate new, accurate information about individuals through inference. This inferred knowledge should be considered personal data under the GDPR when looking at its definition in Article 4(1). This provision contains the word 'any information,' which includes accurate or inaccurate information, collected or generated information. This practice raises important privacy questions that have not been answered yet by authorities and courts.

5. Inconsistent roles and responsibilities of the parties

While the GDPR assigns responsibilities for the ownership and processing of personal data (data controller and data processor), the AIA outlines roles related to the creation and deployment of AI systems (provider, deployer, distributor, and importer). This distinction shifts the burden of risk assessment, creating challenges in aligning overlapping obligations regarding risk evaluation and accountability. Additionally, this divergence complicates coordination and compliance, particularly in informing individuals about data processing and safeguarding their rights. Clear delineation and agreement on these responsibilities will be crucial among the involved parties.

6. Governance Disparities

The AIA's complex governance provisions designate roles for various authorities responsible for monitoring and enforcing its application. This complexity introduces risks of fragmentation and regulatory friction, especially when sector regulators and data protection authorities (DPAs) have overlapping or conflicting responsibilities. The lack of clear procedures for DPA involvement and consistent interpretation among authorities exacerbates these challenges.

A proactive and practical approach to preserve privacy in AI systems

To comply with both the GDPR and the AIA, it is crucial to address the challenges posed by these regulations in both existing and new AI projects. Navigating the distinct requirements of each regulation effectively is essential. While merging the two regulations might seem logical, treating GDPR requirements as a subset of AIA requirements or vice versa is impractical due to their fundamental differences in focus and risk assessment methods. The GDPR imposes requirements on data processing activities, whereas the AIA targets the AI system as a product. Additionally, their risk assessment criteria differ: the GDPR evaluates risks based on specific criteria, while the AIA categorizes risks into unacceptable, high, and limited based on use cases.

A practical approach would be to treat the two regulations separately, listing their requirements for each project. The goal is to create a complementary set of requirements while minimizing redundancy. Where conflicts arise, a risk-based and balanced approach should be employed. These conflicts are likely to be addressed in future court cases, so staying updated on judicial decisions will be crucial for adapting AI project strategies accordingly.

Next, we will clarify the context of the AI system, determine the requirements of both legislations, and implement these requirements into the lifecycle of an AI system. By maintaining clear and separate lists of GDPR and AIA requirements, organizations can ensure comprehensive compliance and mitigate potential legal and operational risks.

1. Understanding the context of AI systems

The first step is to understand the context of the AI system. This involves defining the problem the AI aims to solve, identifying the data used for training and testing, determining the expected outputs, and assessing how those outputs will be used. Organizations must also evaluate the potential risks of flawed or biased outputs to ensure the system is designed ethically and responsibly.

Additionally, it is important to clarify whether the AI system encompasses one or multiple processing purposes. Each processing purpose must meet GDPR requirements for lawfulness and transparency.

2. Determining GDPR and AIA requirements

Organizations must separately evaluate the requirements of the GDPR and the AIA. For the GDPR, this includes assessing the criticality of personal data processing, identifying roles such as data controllers and processors, and ensuring compliance with privacy-by-design principles. For the AIA, compliance depends on the AI system's risk level, ranging from unacceptable risk (prohibited systems) to minimal risk (no obligations).

High-risk AI systems under the AIA must undergo conformity assessments and adhere to strict requirements, such as fundamental rights assessments and transparency measures. By maintaining clear lists of GDPR and AIA obligations, organizations can minimize redundancy and navigate conflicts effectively.

	GDPR	AI Act
Criticality assessment	The criticality level is based on what, how much and for what purpose personal data will be processed, and their impact on the rights and freedoms of individuals.	Under the AIA, use cases are defined by risk levels, being unacceptable risk, high-risk, limited risk and minimal risk. Determining the criticality here is a matter of comparing the use case with the use cases listed in the AIA.
Stakeholder identification	The GDPR lays down roles based on who has ownership of the personal data (the data controller) and who merely processes the data (the data processor). Data controllers have the obligation to follow all requirements laid down on the processing activity, whereas data processors process personal data that complies with instructions of data controllers.	The AIA defines roles based on the party who created and puts the AI system to use. A difference is made between a provider, a deployer, a distributor and an importer. Most of the requirements will be invoked on the creator of the AI system (the provider). The party who puts the AI system to use (the deployer, distributor, importer) will need to make sure that the provider did the necessary to ensure compliance.
Compliance requirements	All processing activities will need to comply with the GDPR requirements. New processing needs to be subject to a privacy-by-design approach. High-risk processing will require an additional DPIA.	The AIA's overall risk-based approach means that, depending on the level of risk of the use case, different requirements apply. (1) unacceptable risk, in which case AI systems are prohibited; (2) high risk, in which case AI systems are subject to extensive requirements, including e.g. fundamental rights assessments and conformity assessments; (3) limited risk, which triggers only transparency requirements; and (4) minimal risk, which does not trigger any obligations. Some of these obligations apply to providers of AI systems, while others apply to the deployers. It is worth noting that for good practice, all AI systems should follow the ethics guidelines by the High-Level Expert Group on Trustworthy AI. These can be used as a starting point upon which AIA requirements can be added.

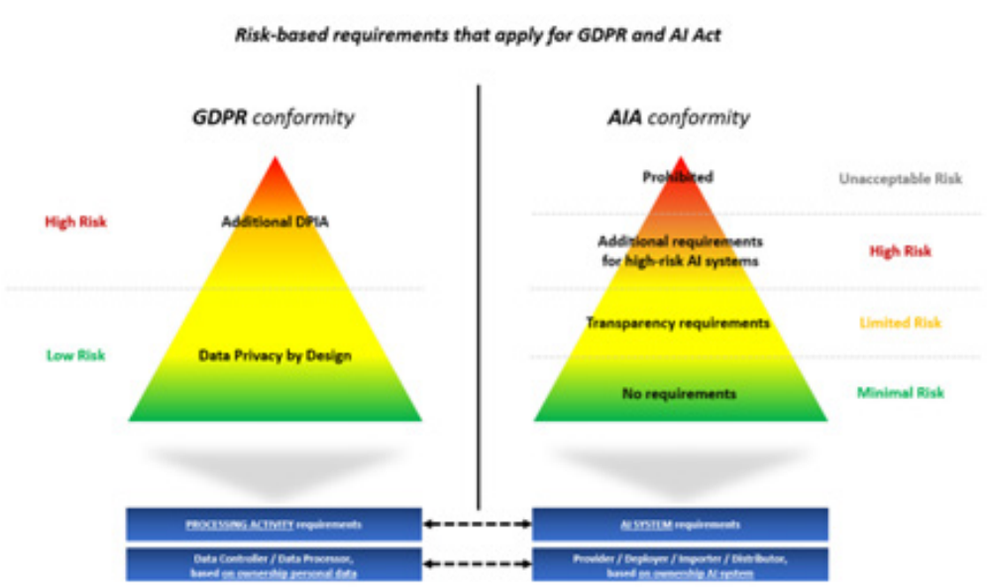

Fig. 1: Diagram of the risk-based requirements that apply for the GDPR and AIA.

3. Implementing requirements in the AI lifecycle

To ensure compliance, GDPR and AIA requirements must be integrated into the AI system's lifecycle, which typically includes design, development, and deployment phases. During the development phase, organizations should consider whether personal data is truly necessary for training the model. If personal data is used, appropriate safeguards must be implemented, such as pseudonymization and encryption. Transparency is critical at every stage. Organizations should clearly communicate how data is used and how the AI system operates.

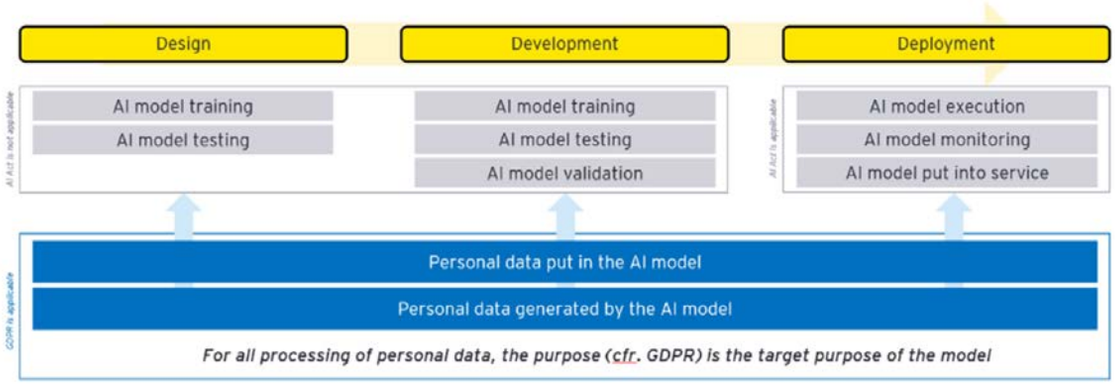

Fig. 2: AI system lifecycle activities and regulatory considerations.

4. Project execution and monitoring the requirements

Once the AI system is operational, organizations must continuously monitor its adherence to GDPR and AIA requirements. Regular audits, reporting mechanisms, and updates to the system's design are essential for maintaining compliance. Organizations should also stay informed about evolving regulations and judicial decisions to adapt their strategies as needed.

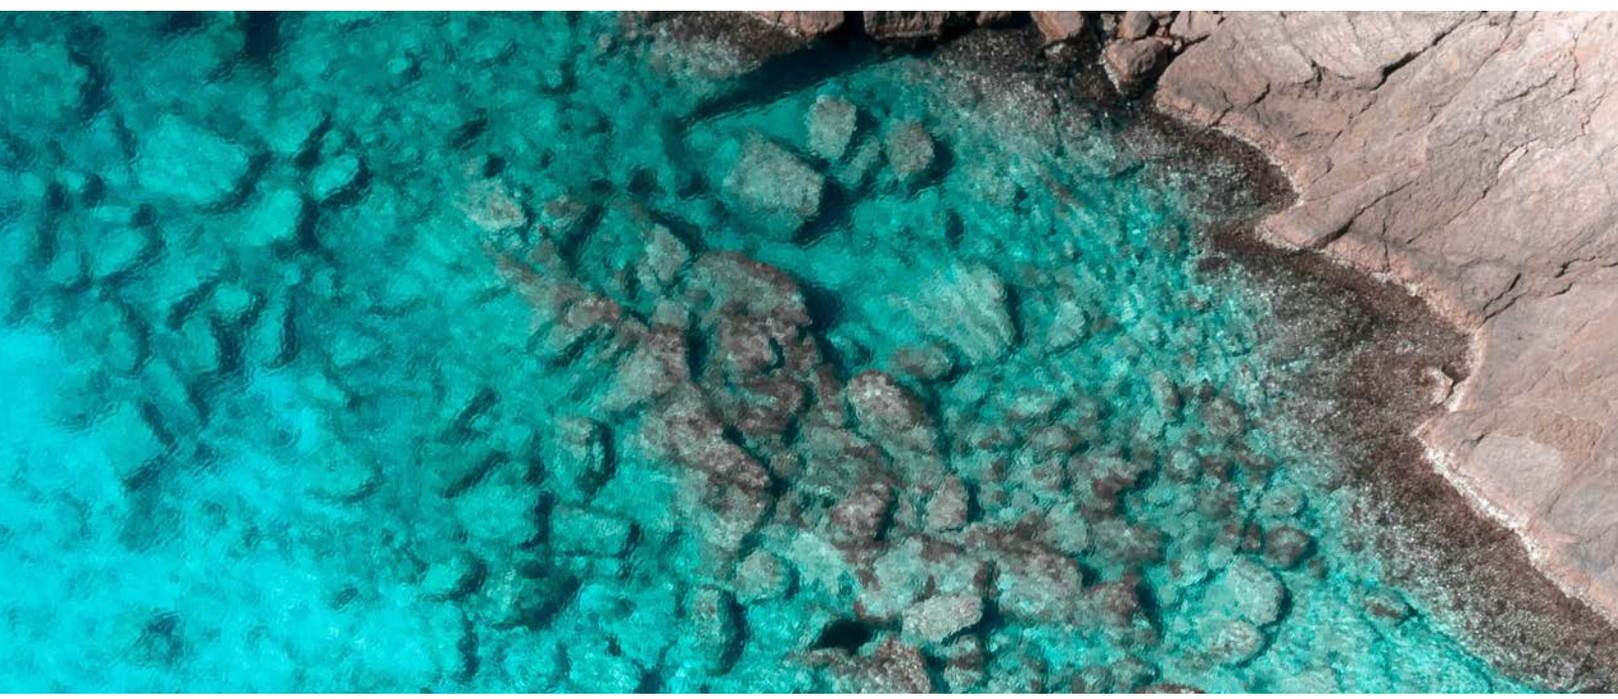

Conclusion

The combined application of the GDPR and the AIA ensures that AI technologies are developed responsibly and ethically. By adopting a structured approach that separates and aligns the requirements of both frameworks, organizations can mitigate legal and operational risks while fostering innovation. As AI continues to evolve, proactive compliance strategies and ongoing monitoring are essential to uphold privacy and protect fundamental rights in the age of AI.

Our team of experts is here to guide and assist you with applying the AI Act and GDPR. EY can help you assess your readiness, define your roadmap to compliance, identify your scope, and perform impact assessments to ensure your company remains secure and compliant.

Contact

Koen Machilsen - koen.machilsen@be.ey.com
EY Consulting Partner

Raf Ganseman - raf.ganseman@be.ey.com
EY Consulting Partner

Yannick Scheelen - yannick.scheelen@be.ey.com
EY Consulting Executive Director

Edwin Haedens - edwin.haedens@be.ey.com
EY Consulting Executive Director

Thomas Diependaele - thomas.diependaele@be.ey.com
EY Consulting Manager