EU POLICY NOTE

#FERMAFOCUS #DIGITLATISATION

October 2024

The EU's Artificial Intelligence Act

Background

On 12 July, the EU AI Act was published in the Official Journal of the European Union[1]. A world-first in terms of horizontal legislation covering uses of AI, the Act sets out to ensure a high level of protection of health, safety and fundamental rights against the potential harmful effects of AI systems.

The EU AI Act is a regulation, which means it will apply to all 27 EU Member States in the same way. At the heart of the Act is a 'risk-based' approach to regulating uses of AI, which will likely become a template for AI law across the rest of the world. In fact, the EU's own AI Office, which has been set up as part of the EU AI Act, has described the Act as having a risk management logic[2].

Additionally, the focus extends to the trustworthiness of AI systems, heavily influenced by the EU product legislation aimed at preventing and mitigating safety risks linked to these products.

It's important to note that the EU considers the AI Act as complementary to existing legislation. As a result, the Act does not include specific liability rules. This important aspect is addressed, however, in the proposed AI Liability Directive (AILD), as well as the existing Product Liability Directive and national Tort Law.

This note walks through some of the major cornerstones in the EU AI Act of relevance to risk managers. It reflects on the *practical* requirements and presents considerations concerning *insurance*.

The Risk-Based Approach

In the risk-based approach of the EU AI Act, a 'pyramid of criticality' classifies uses of AI systems based on the risks they pose to health, safety and fundamental rights: i) low or minimal risk, ii) limited risk, iii) high risk; and iv) unacceptable risk [see diagram on next page].

The use of AI systems that present an unacceptable risk is prohibited. An example would be AI systems used for social scoring, or those employing exploitative techniques to manipulate a person's behaviour, resulting in harm.

For systems classified as minimal or limited risk, such as chatbots or biometric categorisation systems, providers (developer of an AI system) and deployers (did not develop but make use) must disclose that AI has been used.

The bulk of the regulatory requirements apply to high-risk AI systems. These systems must be registered in an EU database before being placed onto the market and must comply with related obligations, including data training and governance, transparence and, risk management systems [Article 9 of the Act].

For general-purpose AI models, the EU distinguishes between those that pose a

[1] Regulation (EU) 2024/1689: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689
[2] See here: https://digital-strategy.ec.europa.eu/en/events/1st-european-ai-office-webinar-risk-management-logic-ai-act-and-related-standards

systemic risk compared to those that do not. The European Commission has the power to classify a general-purpose AI model as posing a systemic risk.

Risk defined: Article 3(2) of the Act defines risk as *'the combination of the probability of an occurrence of harm and the severity of that harm'*. This definition differs to the COSO definition of risk but is importantly not mutually exclusive. AI risks can still be a subset of risks in a broader sense.

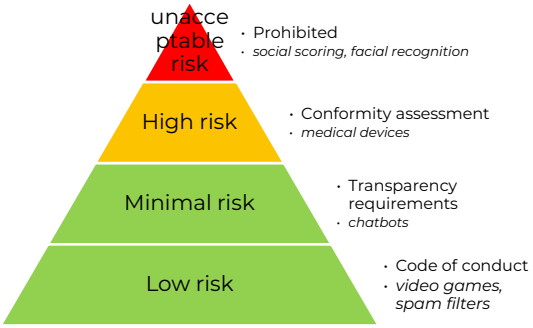

General-purpose AI models
Transparency rules for all models
Risk management requirements in case of systemic risk posed

Key dates

- **12 July 2024 :** EU AI Act published in the Official Journal of the EU.
- **Q1 2025:** AI literacy and training obligations take effect, as do Prohibited AI restrictions
- **Q2 2025:** EU AI Office to publish Codes of Practice relating to General-purpose AI Models
- **Q3 2025:** General-purpose AI Model obligations take effect
- **Q1 2026:** Implementing act and guidelines for High-risk AI Systems
- **Q3 2026:** rules on penalties and remainder of Act comes into effect

Risk management focus

Organisations that develop or use AI systems in their daily operations must implement specific risk management measures. The risk management measures outlined in the EU AI Act will likely be a subset of a more comprehensive approach to risk management at enterprise level. Nevertheless, it is important for enterprises using AI to demonstrate a dedicated focus on managing AI-related risks.

In the language of the legal text, providers and deployers of high-risk AI systems must establish a risk management system, which meets the specifications of Article 9 of the Act.

A deeper look at Article 9

Article 9 of the Act, which is in section 2 of Chapter III entitled 'Requirements for High-Risk AI Systems' sets out the necessary elements for such a Risk Management System [refer to table on next page].

The first clause of Article 9 requires those organisations making use of AI systems to establish, implement, document and maintain a risk management system in relation to high-risk AI systems. This system must be updated throughout the system's lifecycle—for which there is no definition in the text.

The risk management system specified in the text can be viewed, in broad terms, as the steps that the provider or deployer must follow to establish the system [Article 9(2a-d)], along with the actions they may take to mitigate risks.

Three additional points can be made regarding the risks and risk management measures.

1. In the context of high-risk AI systems, risks refer solely to those that can be reasonably mitigated or eliminated through the development or design of the high-risk AI system, or by providing sufficient technical information.

2. When implementing the risk management measures required in Article 9, enterprises need to consider

the effects and possible interactions arising from their combined application.

3. The risk management measures shall strive to find an acceptable level of residual risk both with each hazard and the overall residual risk of the high-risk AI systems.

In addition, the EU requires users of high-risk AI systems to test their systems [Article 9(6-8)]:

- For the purpose of identifying the most appropriate and targeted risk management measures
- To ensure that high-risk AI systems perform consistently for their intended purpose
- To ensure that they comply with Article 9 requirements

The testing is to be done before the AI system is placed on the market or put into service and should be carried out against pre- defined metrics and probabilistic thresholds. These prior defined metrics and probabilistic thresholds are not yet specified, however, and will likely be driven by context.

Lastly, the risk management system should include consideration of whether the AI system is likely to have an adverse impact on persons under the age of 18, and where appropriate, other vulnerable groups [Article 9(9)].

Risk management step & measure table (from Article 9)

REFERENCE	STEP	REFERENCE	MEASURE
ARTICLE 9 (2A)	Identify and analyse the known and the reasonably foreseeable risks that the high-risk AI system can pose to health, safety or fundamental rights when the high-risk AI system is used in accordance with its intended purpose	**Article 9 (5a)**	Eliminate or reduce risks identified and evaluated in as far as technically feasible through adequate design and development of the high-risk AI system
9(2B)	Estimate and evaluate the risks that may emerge when the high-risk AI system is used in accordance with its intended purpose, and under conditions of reasonably foreseeable misuse	**9 (5b)**	Where appropriate, implement adequate mitigation and control measures addressing risks that cannot be eliminated
9(2C)	Evaluate other risks possibly arising, based on the analysis of data gathered from the post-market monitoring system [in other words the collection, documentation and analysis of relevant data once the AI system is on the market]	**9 (5c)**	Provide information in line with transparency requirements [Article 13] and, where appropriate, training to deployers. With a view to eliminating or reducing risks related to the use of the high-risk AI system, due consideration shall be given to the technical knowledge, experience, education, the training to be expected by the deployer, and the presumable context in which the system is intended to be used.
9(2D)	Adopt appropriate and targeted risk management measures designed to address the risks		

Source: FERMA, based on EU AI Act

Further considerations

As with any legal text, the EU AI Act has some terminology that could be considered as open to interpretation. In as much as this causes problems, it also provides some leeway.

Take for example the term '*reasonably foreseeable risks*' [Article 9(2a)] about which there will likely be some discussion or uncertainty over what is reasonably foreseeable, as opposed to unreasonably foreseeable, or even reasonably unforeseeable. Here, each Risk Manager will develop their own method but in the language of likelihood of risks, if a risk is highly unlikely it is still foreseeable.

Another term open to interpretation appears in the context of the need to minimise residual risk to a level that is '*judged to be acceptable*' [Article 9(5)]. The question of acceptable to whom is likely to be raised. Risk Managers may approach this in a variety of ways but documenting the process of the measures and demonstrating how the combination of those measures reduce the risk should be a good starting point.

As a final reflection, there is a mention in the Recitals to the '*state of the art in AI*' as well as incomplete standards or points of reference in some important areas (e.g. probabilistic thresholds, mitigation and control measures). These areas of the text draw out the reality that many will have to approach these measures on a best-effort basis.

Enforcement

According to Article 28 of the Act, EU Member States will establish at least one notifying authority and at least one market surveillance authority. The market surveillance authority will be primarily responsible for the enforcement of the EU AI Act at national level. If an AI system is non-compliant, action can be taken.

Where the market surveillance authority finds that there is either i) non-compliance with the obligations of the EU AI Act – including the Risk Management System requirements, or ii) compliance but nevertheless the high-risk AI system persists in presenting risks to the health, safety or fundamental rights of persons or public interest then enforcement measures can be taken, i.e. the AI system can be prohibited from the market.

On penalties, the text stipulates either up to €35 million or up to 7% of the company's worldwide annual turnover for non-compliance, whichever is higher. In addition, there could be penalties of up to 1%of a company's total worldwide annual turnover for the supply of incorrect, incomplete, or misleading information.

Finally, there are also implications for liability claims linked to AI since there is a proposal for an AI Liability Directive. The logic behind the AI Liability Directive (i.e. reversing the burden of proof from the injured party to the developer) could open up a swathe of litigation against developers or producers of AI systems. This is certainly something for Risk Managers to factor into their considerations.

Practical considerations for Risk Managers

Risk managers must acknowledge the importance of fostering customer trust in an organisation's capacity to develop, deploy, use, and sell AI systems in a responsible and ethical way, while ensuring data is handled responsibly. Achieving this requires the implementation, ongoing review, and documentation of three essential pillars:

1. Development of an AI strategy and transposition into a suitable governance framework, which can be demonstrated by a policy document and end-to-end processes implementation.
2. Implementation of the appropriate technology and investment in the continuous training of employees and partners, as well as providing documentation and guidance for customers.
3. Governance and technology are designed in a way that anticipates audit requirements; and, pursuing a formal

certification is recommended, although not explicitly required by law.

In more detail:

The policy document mentioned above should define, at a high level, the rules and principles for the AI system's lifecycle and how to implement them. Risk Managers can also consider implementing the following elements:

To facilitate maintenance and upkeep of the policy, it is **recommended to follow an internationally recognised ethical 'standard'** such as the Recommendation on the Ethics of Artificial Intelligence, published by UNESCO, to set the groundwork of the principles.

Clearly define the ***scope*** of the policy and the ***roles*** and ***responsibilities***

Consider the scope of the environment in which the AI system operates, checking to what extent existing data controls or governance are adequate or appropriate for the AI tool.

Operationalisation and training are crucial. Do not expect the 'old' organisation to just additionally put the principles into action, and so consider forming a new dedicated entity with close links to Risk Management.

Companies also have to **invest in safe technology implementation**, as well as training. These elements may be helpful as a guide:

Decide on a **technical implementation**. In the case of a Large Language Model (LLM) either install an (open source) locally or protect the business data by access to the LLM via a secure access layer, for example, by leveraging business assistants.

Consider creating an **internal** set of **benchmarks** or tests that measure the performance of the AI system or tool against tasks which are economically valuable to your organisation, and keeping that as a reference point to compare against over time in order to gauge progress, or efficiency gains.

Leverage Retrieval Augmented Generation (RAG): LLMs often lack the specificity and context of specific business data. The lack of business context in LLMs leads to the risk of hallucinations or incorrectly generated information. RAG offers contextualisation of idiosyncratic business data instead of costly training of a dedicated local LLM system or high-risk or even non-compliant use of the business data in LLM prompts.

Ensure users are trained to mitigate the risk of misuse, unethical outcomes, potential biases, inaccuracy, and data and security breaches. In particular:

1. Ensure users are trained to never include personal data or special categories of personal data.

2. Ensure users are using AI systems for business purposes only; do not allow personal use.

3. All uses need to align with the AI policy. Ensure staff do not circumvent by directly using publicly accessible LLMs via the Internet.

4. Pay special attention to training employees to identify situations where the AI tools may operate in ways that are very dissimilar to humans.

5. Identify clearly where data or artefacts are the product of an AI process so they can be differentiated from non-AI outputs.

6. Consider extensive auditing and logging of AI tools, especially along human-to-AI boundaries in order to capture the ways that the tools are being used throughout the organisation.

As mentioned above, one must *anticipate* auditing requirements since the AI Act does not specify an auditing framework for AI. However, adopting some of these elements may assist in being prepared for the eventual audits:

By establishing an **AI Risk Management Framework**, for example the NIST[3], and the required controls, such as Steering Committees, to evaluate potential high-risk use cases. The NIST Framework is unfortunately not an auditable standard, but implementation in accordance with an auditable standard is what is required for answering customer's requests in a fast and cost efficient but still satisfactory way.

It is **recommended to further certify systems** according to an internationally recognised standard. ISO/IEC 42001 is an international standard that specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) within organisations. It is designed for entities providing or using AI-based products or services, ensuring responsible development and use of AI systems.

Illustration of insurance considerations for Risk Managers

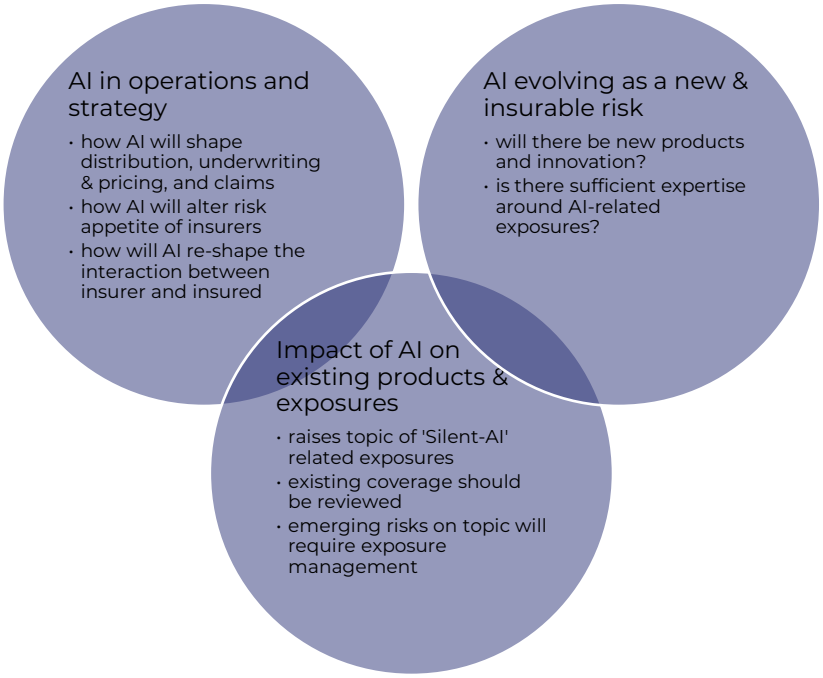

Source: FERMA, based on input from Lloyd's of London

[3] FERMA notes that, currently, there is only a limited sample of 'existing' frameworks and acknowledges that there is work to be done by the EU Standards bodies/EU AI Office or even ENISA, in cooperation with the Risk Management community, in producing a 'standard' or a 'reference'.

Insurance considerations for Risk Managers

Broadly speaking, the impacts of AI on the insurance market are well-documented. In 2021, for instance, EIOPA highlighted that among the benefits AI can bring to insurance such as '*prediction accuracy, automation, new products and services, and cost reductions*', it will also be a challenge for insurers to ensure '*fairness, non-discrimination, transparency and explainability*'.[4]

The diagram above presents one way of categorising the major forces influencing the relationship between AI and insurance. AI will shape the operations and strategies of insurers, from impacting distribution models to forcing them to consider the training or up-skilling needs of staff such as underwriters, claims handlers and so on. It will also impact upon current and future product development and governance.

Less attention, however, has been paid to how the impact of AI on insurers is then felt by corporate risk and insurance managers. For this demographic, it will be telling to see how AI impacts upon existing products and exposures, as well as how it might evolve as a new and insurable risk.

Risk Managers should consider analysing potential 'Silent-AI'—the unknown or unquantified exposures to AI that sits in other insurance policies currently. They should review existing products to see the extent uses of AI in industrial processes (or otherwise) are covered.

Further into the future, Risk Managers might consider evaluating their need for a new type of product related to the way their enterprise uses AI, in line with risk appetite and estimated exposures.

As always, Risk Managers must also stay vigilant to changes in legislation that also have knock-on impacts to their risk transfer strategies concerning AI. For example, as previously mentioned, the AILD might introduce legal requirements that would have far ranging impacts on how Risk Managers might need to calculate their exposures, as well as the appropriate cover they would need to buy for their organisations.

FERMA

The Federation of European Risk Management Associations brings together 23 risk management associations in 22 European countries, representing over 5600 risk managers active in a wide range of organisations. FERMA provides the means of co-ordinating risk management and optimising the impact of these associations outside their national boundaries on a European level.

www.ferma.eu

Contacts
Charles LOW
charles.low@ferma.eu
Typhaine BEAUPERIN
typhaine.beauperin@ferma.eu

[4] EIOPA report on artificial intelligence governance principles; https://www.eiopa.europa.eu/eiopa-publishes-report-artificial-intelligence-governance-principles-2021-06-17_en