Consilio Institute: White Paper

THE EU AI ACT: PRACTICAL STEPS TO PREPARE

Xavier Diokno
Vice President, Solutions & Innovation

Amber Foster
Consultant Legal Counsel

Consilio / ADVANCED LEARNING INSTITUTE

Consilio Institute: White Paper

THE EU AI ACT: PRACTICAL STEPS TO PREPARE

CONTENTS

THE EU AI ACT

On March 21, 2024, the European Parliament announced the final approval of the EU Artificial Intelligence Act (Act). The Act aims to ensure safety and the protection of people's fundamental rights, while still promoting innovation.

Once the final language is validated, the Act is expected to be formally endorsed in June 2024 by the Council of the European Union. It will enter into force 20 days after publication in the Official Journal, but with some elements of the Act coming into force later. In this paper, we give an overview of the Act and set out the practical steps an organisation needs to undertake to prepare and to assess whether the Act will apply to them.

Do not assume that if you are headquartered outside of the EU, the Act will not apply. Even if you find it applies to you unexpectedly, taking the steps outlined here will help you to put in place the systems, risk assessments, and training necessary to implement AI responsibly and safeguard people's rights. Beyond any obligation under the Act, this is positive change and one that will put you in good stead as more regulations come online across the globe.

What Is the EU AI Act and Who Is Affected?

The Act is a set of regulations that govern the placing on the market, putting into service, or use of artificial intelligence systems in the EU. It applies to organizations that provide, deploy, import, or distribute AI systems, which can include organizations outside of the EU if the output of an AI system is used within the EU. The Act describes "AI systems" as being autonomous, machine-based systems that produce output based on the inferences made from data. This includes Generative AI (GenAI) applications that support applications such as ChatGPT.

A RISK-BASED APPROACH

A Brief Overview

The Act will classify AI systems into four risk categories: AI systems posing "unacceptable risks" (e.g., deceptive or manipulative techniques that impair decision-making) will be banned with some limited law enforcement exceptions. Others will be deemed "high risk" (e.g., transportation infrastructures or medical devices), while many AI tools (e.g., chatbots and spam filters) will be deemed lower or minimal risk.

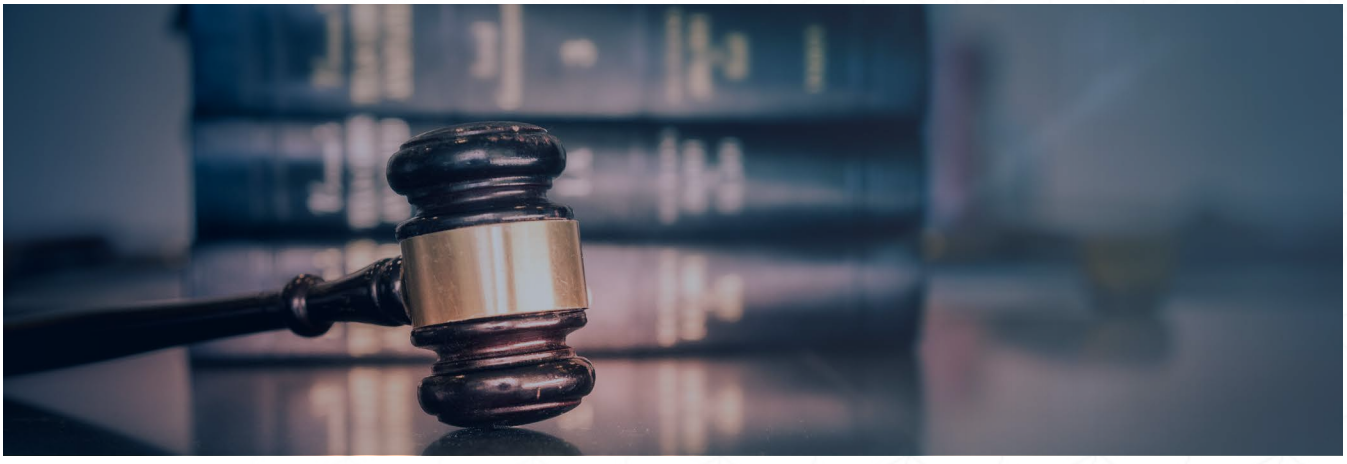

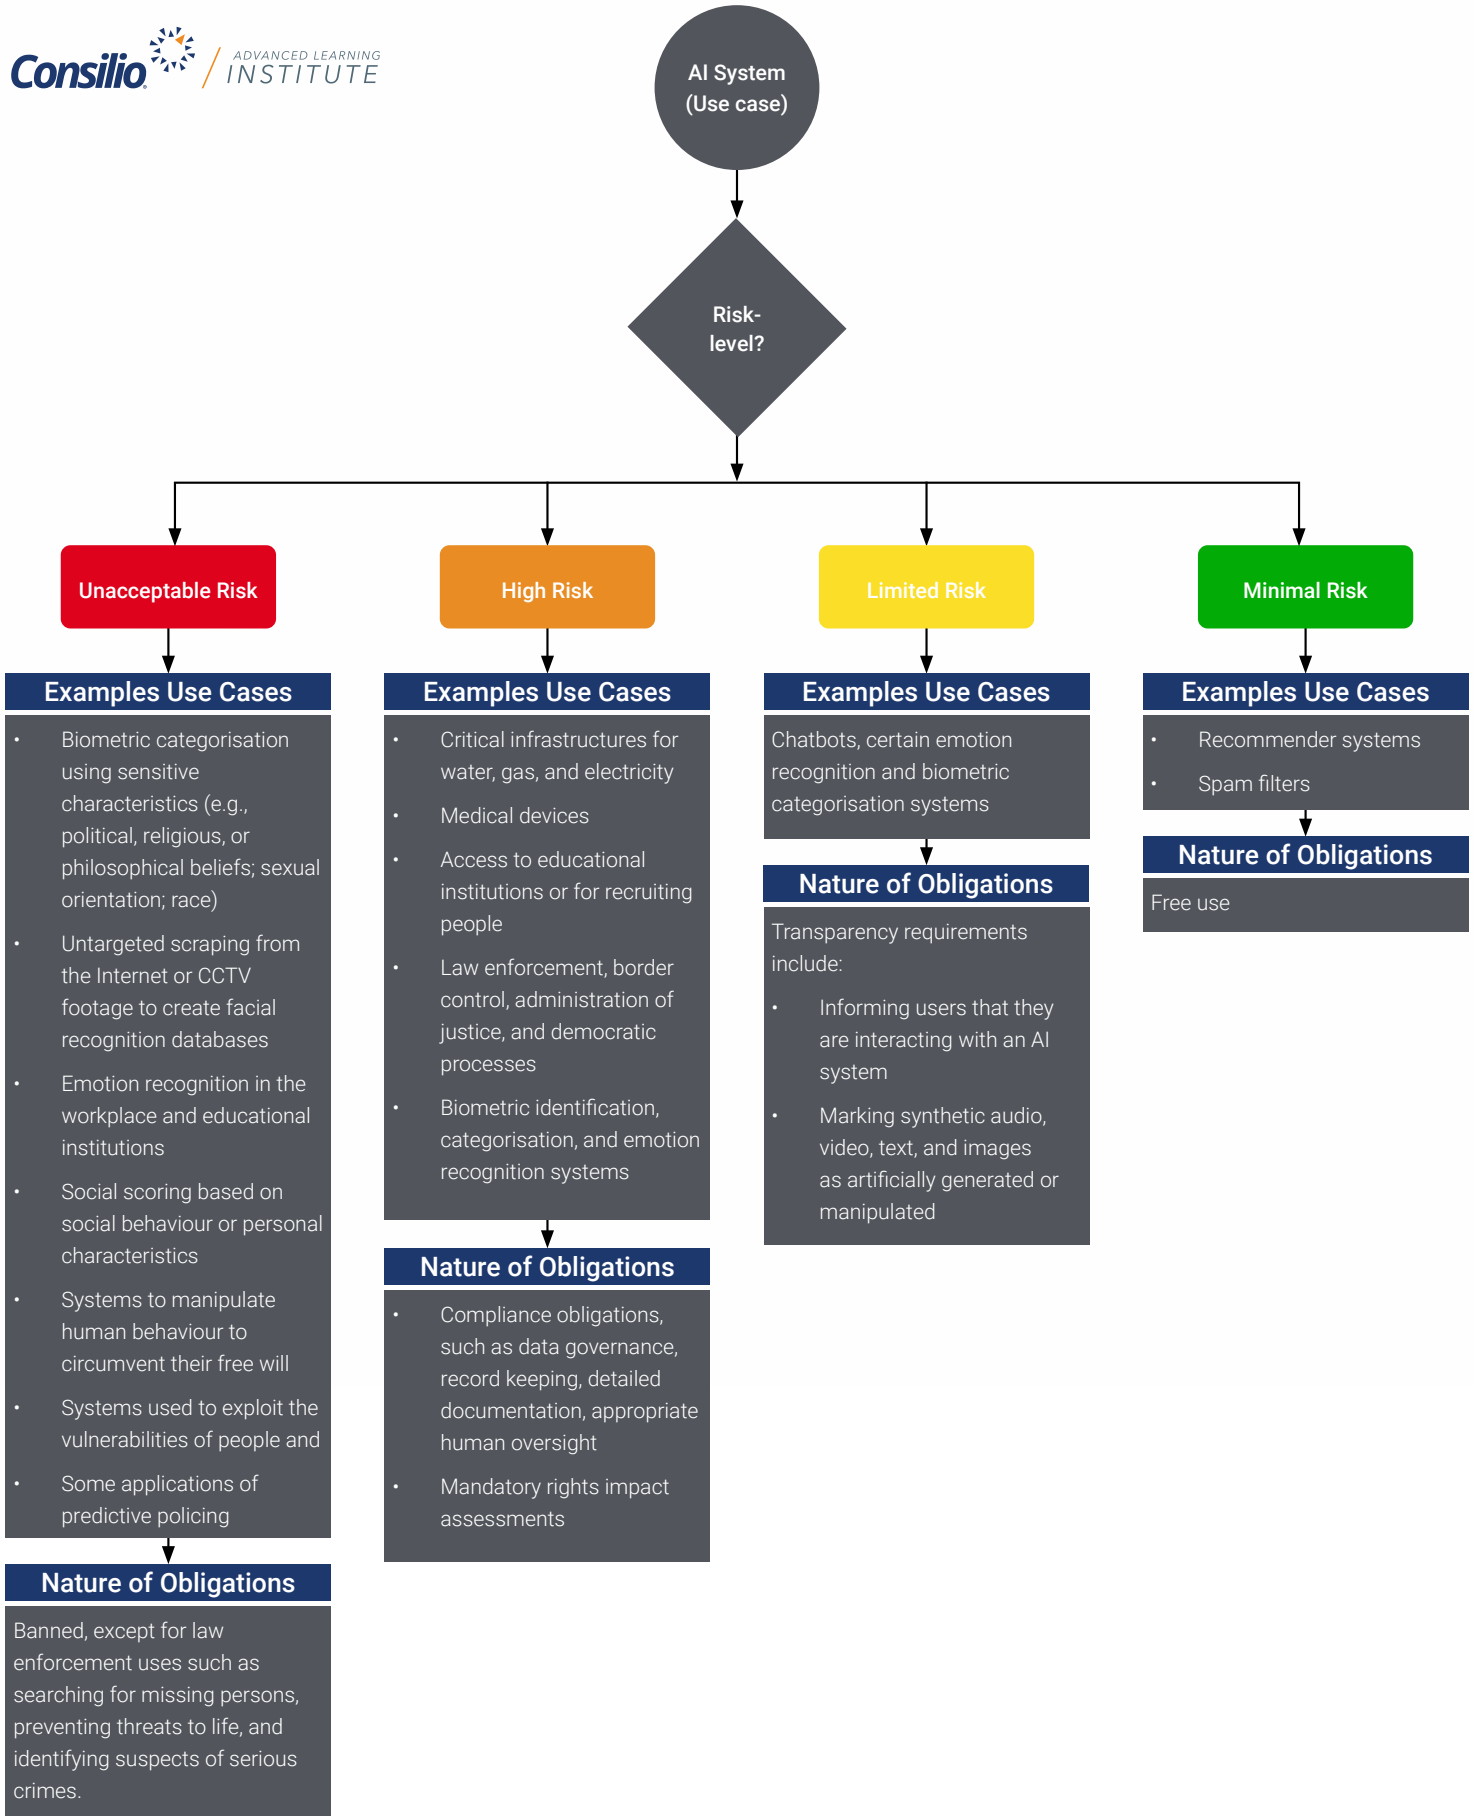
AI System (Use case)
Risk-level?
Unacceptable Risk
Examples Use Cases
Biometric categorisation using sensitive characteristics (e.g., political, religious, or philosophical beliefs; sexual orientation; race)
Untargeted scraping from the Internet or CCTV footage to create facial recognition databases
Emotion recognition in the workplace and educational institutions
Social scoring based on social behaviour or personal characteristics
Systems to manipulate human behaviour to circumvent their free will
Systems used to exploit the vulnerabilities of people and
Some applications of predictive policing
Nature of Obligations
Banned, except for law enforcement uses such as searching for missing persons, preventing threats to life, and identifying suspects of serious crimes.
High Risk
Examples Use Cases
Critical infrastructures for water, gas, and electricity
Medical devices
Access to educational institutions or for recruiting people
Law enforcement, border control, administration of justice, and democratic processes
Biometric identification, categorisation, and emotion recognition systems
Nature of Obligations
Compliance obligations, such as data governance, record keeping, detailed documentation, appropriate human oversight
Mandatory rights impact assessments
Limited Risk
Examples Use Cases
Chatbots, certain emotion recognition and biometric categorisation systems
Nature of Obligations
Transparency requirements include:
Informing users that they are interacting with an AI system
Marking synthetic audio, video, text, and images as artificially generated or manipulated
Minimal Risk
Examples Use Cases
Recommender systems
Spam filters
Nature of Obligations
Free use

GENERATIVE AI

What about Generative AI Systems?

GenAI applications use Large Language Models (LLMs) to understand and generate humanlike responses. Often referred to as "foundational models," these models have been trained on massive datasets and offer a broad range of capabilities. The Act refers to these models as General Purpose AI (GPAI) models. Applications that use GPAI models are referred to as GPAI systems. Some popular GPAI systems include OpenAI's ChatGPT, Microsoft Copilot, and Google Bard. These applications have a broad range of capabilities, including acting as chat assistants or analyzing, creating, or summarizing content.

Providers of GPAI models have obligations around transparency (documentation), copyright law compliance, and summarizing the data used to train the model. Some examples include creating technical documentation that describe the training, testing, and evaluation process, providing documentation to downstream organizations that use these models, and a detailed summary about the content used for training the model. Providers of free and open license GPAI models only need to comply with copyright law and publish the training data summary unless the model poses a systemic risk.

Some GPAIs will be seen as having systemic risk if they have the potential to: have a high impact (e.g., major accidents, serious public health consequences); be misused for cyberattacks; harm individuals through bias, discrimination, or disinformation; or if the Commission has identified them as having a systemic risk for another reason. Providers of these systems must comply with additional codes of practice, such as performing standardized model evaluations, assessing and mitigating risks, tracking and reporting serious incidents, and ensuring adequate cybersecurity protections.

GETTING READY

5 Key Questions for Your Organization

1. **Will it apply to my organisation even if we are not in the EU?**

   As the Act has a broad scope, it could potentially apply to your organisation even if it is not in the EU.

   Developers, deployers, importers, and distributors of AI systems will be covered if their **output occurs within the EU**. It is essential that users of AI systems work out the role they are playing in the supply chain (see below) and audit their use of the system.

   Even if the Act does not apply to your organisation, it will likely be seen as a benchmark for excellence and provide companies with a model for analysing AI risks and identifying mitigation strategies.

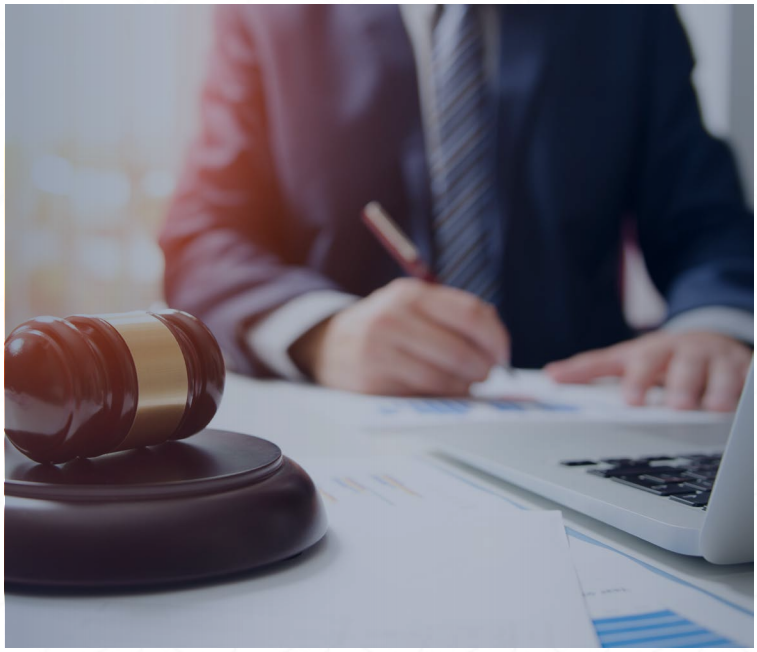

Where does my organization sit in the AI supply chain?

Role	Description
Provider/Manufacturer	Develops or intends to put an AI system on the EU market
Importer*	Places an AI system on the market under the trademark of a person/entity established outside of EU
Deployer*	Uses AI tools
Distributor*	Distributes AI systems and models to the EU

*These roles can become Providers if they market an AI system under their own trademark (despite being developed by a third party), if they materially change the AI, or if the purpose of the AI system is modified. This means that customization of AI systems by an organisation may mean they inherit the more onerous obligations of a Provider.

2. **I am unsure which of our IT systems are "AI systems" and which are not – what exactly does the Act apply to?**

   There are numerous philosophical debates on what AI is and isn't. For our purposes, it is a:

   > . . . machine-based system 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.

   Generally, these are computer applications that use models to make inferences or predictions from data. These predictions are then used to provide an answer or response to a question or instruction. For example, streaming services often use AI systems to analyze customer viewing data. Based on the inferences made from this analysis, the system can then predict or recommend shows that the customer will likely enjoy.

   To start with, we recommend auditing the IT tools used across the business and labeling them. Even the simplest technical tools seem to be "powered by AI," and you may be surprised by what is being used in your company. Training your procurement teams to recognize when AI is integrated into a product is key so they can ask the right questions.

3. **How do we start to build our compliance system?**

- Audit your IT systems and work out which systems fall under the "AI" definition.
- Assess each AI system against the four risk categories.
- For each AI tool you use, work out your role (developer, deployer, importer, and distributor).
- Understand where the output occurs – EU or outside.
- Once you have this information, you can map out your obligations depending on your role as developer, deployer, importer, or distributor.

4. **If we will be a "deployer" of AI, what do we need to ask our tech suppliers?**

   Many organizations will be in the deployer category. Below are some targeted questions to ask suppliers:

- How is your system classified under the EU AI Act, and what specific compliance measures have you implemented?
- Have you conducted a risk assessment for the system?
- How does the system manage data governance and compliance with EU GDPR and other privacy rules?
- Please provide a comprehensive overview of the system's data sources, algorithms, decision-making processes, and performance metrics.
- In the decision-making process, where should human intervention and oversight be?
- How have you addressed issues of bias?
- What steps have you taken to mitigate against bias?
- Are there any known limitations in the data sets used to train the model?
- Have you signed up for any ethical guidelines or frameworks?
- How will you ensure compliance with the EU AI Act on an on-going basis?

5. **How should we set up our governance framework?**

Not one size fits all, and we have seen different examples of governance ranging from strict ethical AI committees to organizations treating it like any other business risk.

The basics of setting up an AI governance framework include:

- **Defining the objectives of your framework.** This may include the scope, the AI systems you currently use and their applications, the regulations that you need to comply with, and the stakeholders involved (e.g., executive leadership, IT, AI system administrators, etc.).
- **Engaging your people.** Responsible use of AI is more than just technical, privacy or Infosec concerns. People throughout organizations see the amazing productivity benefits of AI yet are worried about its impact on their job roles, sustainability and issues of bias/ unfairness. We recommend a holistic approach to governance ensuring that HR, Diversity and Inclusion, employee representatives and sustainability teams are involved alongside risk professionals and feel able to call out issues of bias, unfairness and, impact on jobs.

- **Conducting a risk assessment to identify potential areas of AI risk.** Involves identifying the systems that incorporate AI and for each system, gathering system information, its intended uses, any potential adverse impacts, data requirements, and a summary.
- **Defining guardrails for AI use, including what uses are off-limits within the organization.** Guardrails can be either technical or those implemented within the AI system or policies that are enforced by an organization. Example technical guardrails are validation tests performed on AI results, mechanisms that incorporate user feedback, and business rules that define AI behavior. Example policy guardrails may address what data can be used with AI systems, how systems can be designed to comply with data privacy and security requirements, and the intellectual property rights on AI-generated content.
- **Establishing structure and processes, including creating committees, assigning ownership, management of policies, and communication.** The people and tasks that will be responsible for developing, managing, and maintaining your AI governance. This includes identifying the stakeholders involved, members of the committee, member responsibilities, and forms of communication.

KEY TAKEAWAYS

In this paper, we have given you a high-level overview and top-line action plan to start planning for compliance with the Act. Organizations will need to understand where their use of AI falls within the Act's risk classification along with their role within the supply chain. Once an organization determines their risk level and their role, they can determine their compliance obligations, which may include developing an AI governance framework.

ABOUT THE AUTHOR

Xavier Diokno has a bachelor's degree in computer science from Southern Illinois University, a master's degree in computer science from the University of Illinois at Chicago and a juris doctor degree from DePaul University College of Law. He is licensed to practice in the state of Illinois and the United States Patent and Trademark Office.

Prior to becoming an attorney, Xavier worked in the information technology industry for ten years in database administration and software development. For more than a decade, Xavier was part of Consilio's Data Analytics group, where he oversaw the team's tripling in size, as well as numerous large-scale projects involving Technology-Assisted Review, Immediate Case Assessments™, and novel analytics research. Xavier now applies his technical and legal experience to oversee Consilio's Innovation initiatives, including researching new technologies like artificial intelligence and developing their application to legal services.

Xavier Diokno
Vice President,
Solutions and Innovation
m +1.312.638.3130
e xdiokno@consilio.com
consilio.com

ABOUT THE AUTHOR

As an experienced legal consultant, she currently provides in-house legal services to the Lawyers On Demand UK head office and maintains a client base of small entrepreneurial businesses. Amber combines her legal practice with a non-executive portfolio, giving her a deep understanding of strategy, governance and how boards can drive change and innovation.

Amber is passionate about leveraging AI tools to simplify complex legal processes for in-house legal teams and their client colleagues. Her extensive career includes roles as General Counsel and Head of Government Affairs for QVC UK, and as Senior Legal Counsel at Coca-Cola Enterprises, GB. In these positions, she managed a broad range of issues including retail, broadcasting, competition law, and marketing/advertising law.

Amber is a frequent speaker, addressing in-house legal teams and businesses on the practical applications of AI in legal and business settings.

Amber Foster
Consultant, Lawyers On Demand
e amberfoster@lodlaw.com
consilio.com