European Union Artificial Intelligence Act

December 13, 2023

Update: The Act was passed by the European Council and European Parliament on December 9, 2023, thus establishing the world's first regulatory framework to ensure the safety, legality, trustworthiness, and respect for fundamental rights within AI systems.

On June 14, 2023, the European Parliament passed its version of the Artificial Intelligence Act (the "Act"). The Act is now under review by negotiators from the European Union's ("EU") three bodies - the European Commission, Council, and Parliament to reconcile different versions of the Act and finalize implementing language. The European Commission is planning to push for approval of the finalized Act by the end of 2023. When enacted, the Act will be the most comprehensive regulations affecting artificial intelligence ("AI") systems to date.

The Act is the world's first major set of comprehensive rules regulating AI technology.

The Act passed by the European Parliament would restrict what the EU believes to be AI's riskiest uses, such as facial recognition software and chatbots like ChatGPT. The EU's stated goal is to ensure better conditions for the development and use of innovative technologies. The EU recognizes AI's benefits in the healthcare, transportation, manufacturing, and energy sectors but hopes to promulgate regulations that curb potential excesses and violations of EU fundamental rights.

The Act is expansive and would govern any entity providing a service that uses AI. This includes services that produce content, predictions, recommendations, or decisions.

The Act contemplates a risk-based, tiered system for regulating AI.

One of the Act's most important aspects is its attempt to categorize and regulate AI systems based on potential risk. If the Act becomes law, firms using and/or producing AI systems will need to know which category their systems fall under to avoid running afoul of the new regulations. The Act divides AI systems into three distinct categories (unacceptable risk, high risk, and low or minimal risk), which are treated differently by the Act.

Unacceptable Risk

AI systems posing an unacceptable risk are those that are considered a threat to people and are banned under the Act. These systems include cognitive behavioral manipulation, social scoring, and real-time biometric identification systems such as facial recognition. Some exceptions will be allowed. For example, remote biometric identification systems may be allowed when identification occurs after a significant delay so the information can be used to prosecute serious crimes after court approval. EU member states are currently lobbying EU officials to increase exceptions for national security and law enforcement applications of AI systems.

High Risk

High risk systems include those systems that are used to operate critical infrastructure or determine access to legal institutions, public services, and government benefits. These systems are labeled high risk because of their potential to negatively affect the safety and fundamental rights of the EU and its residents. High risk systems will be permitted in the EU subject to certain compliance requirements. High risk systems are further subdivided into two main categories. The first, for example, includes systems that are intended for use as a safety component, medical devices. The second includes other applications that may implicate fundamental rights but do not presently have a separate third-party assessment process already established in law or industry.

The Act proposes the following risk management efforts for high risk systems:

1. Implementation and maintenance of risk management systems for high risk applications;
2. Training, validating, and testing data sets;
3. Up-to-date technical documentation that demonstrates compliance with the Act;
4. Automatic logging of events while an application is operating;
5. Transparency that allows users to interpret AI outputs and use an application appropriately;
6. Human oversight; and
7. Appropriate levels of accuracy, robustness, and cybersecurity.

Obligations for providers of high risk applications include:

1. Ensure compliance with the AI Act;
2. Quality management systems;
3. Providing technical documentation;
4. Keep automatic logs;
5. Ensure high risk applications undergo conformity assessments where applicable;
6. Comply with registration obligations;

7. Take corrective actions where necessary;
8. Inform competent authorities of non-compliance and corrective actions taken, where applicable;
9. Mark high risk applications with required marking to indicate conformity;
10. Demonstrate compliance upon request from competent authorities; and
11. Appoint an authorized representative in the European Union.

Manufacturers, importers, and distributors have further obligations under the Act, but more details remain to be determined during current negotiations over the language of the Act. The Act would require the European Commission to prepare guidance on when harm would trigger the high risk categorization six months before the law takes effect. Providers of impacted AI systems listed as likely to be high risk who do not think that this harm threshold has been met are expected to notify the national supervisory authority accordingly.

Low or Minimal Risk

The Act requires limited risk AI systems to comply with minimal transparency requirements that would allow users to make informed decisions. After interacting with the applications, the user can then decide whether he or she wants to continue using it. Users should be made aware when they are interacting with AI. This includes AI systems that generate or manipulate image, audio, or video content. This would apply to AI systems producing deepfakes. Compared to high risk systems, low or minimal risk systems would be largely unregulated under the current Act.

AI system creators would be required to conduct risk assessments before introducing their AI systems to the public - similar to drug approval processes.

Generative AI systems would be subject to unprecedented transparency regulations under the terms of the Act.

Generative AI software such as ChatGPT will be required to comply with several transparency requirements including:

1. Disclosing that the content was generated by an AI system;
2. Designing the model to prevent it from generating illegal content; and
3. Publishing summaries of copyrighted data used for the system's training.

Providers of foundation models are subject to obligations to undertake risk assessments and mitigate reasonably foreseeable risks; and to establish appropriate data governance measures, obligations relating to the design of the foundation model (including from an environmental impact perspective) and an obligation to register the foundation model in an EU database.

U.S. and other non-EU entities should pay close attention to the regulations proposed in the Act.

Similar to the EU's General Data Protection Regulation, the AI Act will apply extraterritorially to providers putting AI systems on the EU market or employing their use in the EU. This applies to entities regardless of whether they are established in the EU or elsewhere. Importantly, the Act does not have retrospective effect except to systems when there is a significant change to their design or intended purpose after the Act's effective date.

One of the most likely impacts on United States ("US") entities is that compliance with European laws often impacts product and market offerings in the US. For example, if ChatGPT is required to watermark its word product in most of Europe, it is likely to do the same in all markets. For this reason, tech industry leaders in the United States have suggested the Act would be prohibitively difficult to comply with.

The clearest compliance issue involves generative AI that is trained on immense amounts of data. The Act would require detailed summaries of training methods that use copyrighted materials. Most generative AI systems currently have no means of tracking the data used to produce every response. The Act's goal is to ensure copyrights are respected by AI generators, but the Act's language risks spurring litigation in cases when it is unclear whether a particular creator's work was actually accessed by the generative system. This problem is further implicated by the differences between American and European copyright systems. In the EU, there is no copyright registry and any creator automatically has a copyright the moment they create their work. The US, on the other hand, has a systematized copyright system and every work is not eligible for protection. For example, the U.S. Copyright Office released a statement on March 16, 2023, indicating that AI-generated works are not eligible for copyright. Reconciling the two systems poses a serious challenge to US entities and European regulators moving forward.

Failure to comply with the Act could result in multi-billion dollar penalties.

The Act levies substantial fines for entities who fail to comply. Those deemed to have violated the prohibited uses provision would incur a fine up to forty million euros or seven percent of total global revenue for the preceding financial year - whichever is greater. Violations of human-rights laws or any type of discrimination perpetrated by AI would incur fines up to twenty million euros or four percent of global revenue for the preceding year. Other noncompliance offenses, including from foundational models such as generative AI, are subject to fines up to ten million euros or two percent of global revenue for the preceding year. An entity found to be supplying false, incomplete, or misleading information to regulators could be fined up to five million euros or one percent of global revenue for the preceding year.

The compliance penalties expected to be included in the final version of the Act, signaling the EU's intention to make AI regulation a key priority in the coming years.

Enforcement actions would be brought by the EU's individual member states.

Entities doing business in the EU have several years to bring their business practices into compliance with the new regulations.

In short, the EU appears determined to introduce heavy regulation in the AI space. It is unclear how far-reaching this regulation will be but the proposal by the European Parliament would be the most sweeping regulation of AI to date. EU officials expect the Act to adapt to rapidly changing technologies and do not anticipate altering the current framework. The vagueness in many parts of the Act is intentional and would give the EU government significant leeway in interpreting and enforcing the Act.

Moving forward, entities who employ AI systems in their businesses will also want to pay close attention to the Biden administration's "Blueprint for an AI Bill of Rights," which focuses on privacy standards and testing before AI systems are made publicly available. China has also published a set of draft rules that would require chatbot makers to adhere to state censorship laws.

The Act is nearing the end of the trilogue phase of the EU's legislative procedures. The trilogue involves informal negotiations between the European Commission, the European Parliament, and the European Council to reconcile differences in the versions of the AI Act passed by each body. After formal adoption, the legislation will be published with implementation dates and other key details. Full implementation is not expected until 2025.

For more information, contact Dale Werts, or your regular Lathrop GPM contact.

Pierce Rose, a law clerk with Lathrop GPM LLP in summer 2023, assisted with the research and drafting of this article.