GUNDERSON DETTMER

Client Insight: Demystifying the EU AI Act

Insights
November 12, 2024

On August 1, 2024, the landmark comprehensive artificial intelligence ("AI") law – the European Union Artificial Intelligence Act (the "EU AI Act" or the "Act") – entered into force. While most obligations under the Act will not become effective until August 2026, AI systems deemed to pose "unacceptable risk" will be prohibited starting February 2, 2025, and certain obligations, including those for providers of general-purpose AI models, will go into effect August 2, 2025.

The Act has a broad extraterritorial reach, applying to providers and deployers of AI systems, regardless of their location, if the AI system is put onto the EU market or AI system output is used in the EU. This means that many U.S. companies will be subject to the Act even if they have no establishment within the EU. In addition to imposing requirements on AI systems based on their level of risk (with most obligations falling on AI systems considered "high risk"), the Act outright bans certain AI systems considered to pose unacceptable risks to individuals. Violations of the Act will be subject to very significant financial penalties, exceeding the already high fines under the EU General Data Protection Regulation ("GDPR"). Given the extraterritorial reach and the potential for large fines, the EU AI Act is expected to set a global standard for regulation of AI, and performing an early assessment of your organization's level of risk and compliance burdens is strongly recommended. Your Gunderson Dettmer team can help evaluate your company's compliance burden under the Act.

Who Is Subject to the Act?

The EU AI Act applies to providers, deployers, importers, and distributors of AI systems or general-purpose AI models ("GPAI models"), as well as product

manufacturers offering AI as part of their products. Under the Act:

- "**Providers**" are developers of AI systems or GPAI models, or a natural or legal person, public authority, agency or other body that has an AI system or GPAI model developed and places it or puts it into service on the EU market under their own name or trademark. This is generally understood as the vendor or service provider of an AI system.
- "**Deployers**" are users of an AI system. This can be a company purchasing an AI system for use in its organization.
- "**Importers**" are persons that place on the EU market an AI system that bears the name or trademark of a person established outside the EU.
- "**Distributors**" are persons in the supply chain (other than the provider or the importer) that make an AI system available on the EU market.

What Is an "AI system" and Where Does "General Purpose AI Model" Fit In?

The EU AI Act governs AI systems and GPAI models, both of which are broadly defined. An AI system is "a machine-based system designed to operate with varying levels of autonomy, 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." From a practical perspective, tools using deep learning, reinforcement learning, machine learning and natural language processing will generally be AI systems under the Act. The Act **exempts** certain AI systems, such as those used solely for scientific research and development, AI development and testing outside of real-world conditions, as well as AI systems for military, defense, and national security purposes. The Act also **does not apply** to AI systems released under free and open-source licenses, unless they fall into prohibited or high risk categories (*discussed below*).

A GPAI model is not an AI system on its own but can be a critical component of one, and is defined as "an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are released on the market." From a practical perspective, a GPAI model is an AI model with the key characteristics of generality and capability to competently perform a wide range of

distinct tasks. GPAI models may be placed on the market in various methods (e.g., libraries, application programming interfaces ("APIs"), direct download, or physical copies), but require additional components (e.g., a user interface) to be considered an "AI system" under the Act. GPAI models typically have at least a billion parameters and are trained on very large datasets, and include large language models (LLMs) driving many popular AI chat and content generation applications.

Does the Act Apply to Entities Outside of the EU?

Entities located outside of the EU can be subject to the law and its significant penalties. For example, the Act applies to providers who offer AI systems or GPAI models in the EU market, *regardless of where the provider is located.* The Act also applies to deployers and providers located outside of the EU that operate AI systems which produce outputs used in the EU. Practically, this means that the Act casts a wide net of applicability, and many U.S. entities will be subject to the Act even if they have no establishment in the EU and no plans to further expand into European countries. As a result, the Act is expected to set a global standard for compliance as a "highest common denominator," since many entities will find it challenging to deploy different obligations in different jurisdictions with respect to the same product or service.

What Obligations Do Regulated Entities Have Under the Act?

Obligations under the EU AI Act vary depending on the type of regulated entity (e.g., whether you are a provider or a deployer) and the level of risk associated with the AI system.

AI Systems

The Act divides AI systems into four risk levels: (1) unacceptable risk, (2) high risk, (3) limited risk, and (4) minimal risk, and imposes compliance obligations on the first three risk levels (as summarized in the table below), with no obligations on minimal risk AI systems.

AI System Level of Risk	Description	Examples	Treatment
Prohibited Risk	This category includes AI systems that pose unacceptable risks to individuals. This includes AI systems that: • Deploy subliminal, manipulative, or deceptive techniques to materially distort the behavior of individuals by impairing their	Examples include AI systems used in the following contexts: • Facial recognition systems in public spaces; • Social ranking systems that classify people based on behavior, socio-economic status, or personal characteristics;	**Prohibited:** Prohibited risk AI systems are expressly banned.

	ability to make an informed decision; • Exploit vulnerabilities due to age, disability, or social or economic situation to materially distort their behavior in a manner that causes that person (or another person) significant harm; • Conduct social scoring that leads to certain detrimental or unfavorable treatment of the individual; • Assess or predict the risk of an individual committing a criminal offense based on profiling their personality traits and characteristics; • Create or expand facial recognition databases through untargeted scraping of facial images from the internet or CCTV footage; or • Conduct real-time remote biometric identification in publicly accessible spaces for law enforcement purposes, subject to exceptions.	• Voice-activated toys that encourage dangerous behavior in children; or • Essential private and public services, such as automated welfare benefit systems or private sector credit scoring systems.	
High Risk	This category includes AI systems that could be expected to pose significant threats to individuals' health, safety, or fundamental rights, including AI systems in the following areas: • Biometrics (implementation of biometric ID systems in private sector contexts and/or not in real time requires authorization given by a judge or independent authority); • Educational and vocational training; • Employment, workers' management and access to self-employment; • Access to and enjoyment of essential private and public services and benefits; • Law enforcement; • Migration, asylum, and border control management; or • Administration of justice and democratic processes.	Examples include AI systems used in the following contexts: • Critical infrastructure (including digital infrastructure), such as transport, that could put the safety of citizens at risk; • Law enforcement, such as automated risk scoring for bail, deepfake law enforcement detection software, or "pre-crime" detection; • Immigration, such as verification of travel documents or visa processing; • Administration of justice and democratic processes, such as automated sentencing assistance; • Education, such as automated scoring of exams that determine access to education; or • Employment (including recruitment), such as automated hiring or resume sorting software.	**Regulated:** High risk AI systems must adhere to strict compliance requirements, including thorough documentation, high levels of data accuracy, and transparency to ensure traceability and accountability. Providers of high risk AI systems are subject to a number of requirements, including: • Human oversight requirements; • Requirements to design, develop and document the high risk AI systems in a manner to achieve an "appropriate level" of accuracy, robustness, and cybersecurity (including, e.g., automatic recording of events over the lifetime of the system); • Implementation of similar compliance requirements for downstream deployers; • Implementation of data governance procedures, ensuring that training, validation and testing data sets are relevant, sufficiently representative and, to the best extent possible, free of errors and complete according to the intended purpose; • Impact assessment that evaluate impact on fundamental rights; and • Registration with EU authorities, unless the provider can demonstrate that the AI system does not pose a

			significant risk to fundamental rights, public safety, or public health.
Limited Risk	AI systems that pose a risk of impersonation or deception.	Examples include: • Chatbots; or • AI systems that make it possible to manipulate images, sound, or videos (i.e. deepfakes).	**Transparency Obligations:** Providers must ensure users are informed that they are interacting with an AI System. Providers of limited risk AI systems are subject to requirements, including: • Transparency requirements, such as marking AI-generated images, sounds, or video as artificially generated or manipulated; • Compliance with EU copyright law; and • The publication of training data summaries.
Minimal Risk	AI systems that do not fall under the above-mentioned categories that post little to no risk to individuals.	Examples include: • Spam filters; • AI-enabled video games; and • Inventory-management systems.	**No Mandatory Obligations:** Providers have no explicit restrictions or obligations. However, it is recommended to follow basic principles like human oversight, non-discrimination, and fairness.

GPAI Models

The EU AI Act divides GPAI models into "normal" GPAI models and GPAI models with systemic risks, and obligations only apply once the GPAI model is released on the market. "Normal" GPAI models that are released under a free and open-source license are exempted from most of the obligations under the Act. The distinction between these risk categories under the Act is as described below:

GPAI Models	Description	Examples	Treatment
GPAI Models with Systemic Risks	GPAI models are classified as "GPAI models with systemic risk" if they have: • high impact capabilities evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks; or • the capability or impact equivalent to those set out above, based on the criteria set out by the Act (Annex XIII) and based on a decision of the EU Commission, ex officio or following a qualified alert from the scientific panel.	Examples include: • LLMs that are deemed to have systemic risks. Subject to amendments by the EU Commission, current GPAI models that are trained using a total computing power greater than 10^{25} FLOPs* are considered to have systemic risks. *FLOPs (floating floating-point operations per second) is a unit of measurement that	In addition to the obligations listed above for "normal" GPAI models, providers of GPAI models with systemic risks are subject to additional requirements, including: • Model Evaluation – Providers must perform model evaluation in accordance with standardized protocols and tools reflecting the state of the art, including conducting and documenting adversarial testing of the model with a view to identifying and mitigating systemic risks; • Assessment and Mitigation – Providers must assess and mitigate possible systemic risks at the EU level, including their sources that may stem from the development, the placing on the market, or the use of general-purpose AI models with systemic risk; • Documentation and Reporting – Providers must document and report, without undue delay, relevant information about serious incidents

		quantifies the computing power of a computer or processor.	and possible corrective measures to address them; and • Security – Providers must ensure an adequate level of cybersecurity protection for the GPAI models with systemic risk and the physical infrastructure of the model. Further, pending the EU AI Office's forthcoming publication of the "Code of Practice for General-Purpose AI," providers of general-purpose AI models should expect additional requirements where use of models carries "systemic risk."
"Normal" GPAI Models	GPAI models that display significant generality and that are capable of competently performing a wide range of distinctive tasks upon commercial release (e.g., capable of a variety of downstream applications).	Examples include: • Large-scale generative AI models are a common example of a GPAI model, given that they allow for flexible generation of content in the form of text, audio, images or video that can readily accommodate a wide range of distinctive tasks, including OpenAI's ChatGPT.	Providers of "normal" GPAI models are subject to requirements, including: • Documentation – providers must keep detailed records of the AI's development and testing, and provide this information to other companies who want to use it; • Transparency – Providers must mark artificially created or manipulated content as such; • Copyright Compliance – Providers must put a policy in place to comply with EU laws on copyright and related rights; • Training Data Summary – Providers must prepare and make publicly available a detailed summary of the model's training content; and • Cooperation – Providers must cooperate with the EU Commission and national authorities.
	Free and Open-Source "Normal" GPAI models that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available.	Examples include: • LLMs that are free and open-source including Google's LaMDA and Databrick's Dolly.	Providers of free and open-source "normal" GPAI models are only subject to two requirements: • Copyright Compliance – Providers must put a policy in place to comply with EU laws on copyright and related rights; • Training Data Summary – Providers must prepare and make publicly available a detailed summary of the model's training content.

What Are the Key Compliance Deadlines Under the Act?

Compliance deadlines under the Act are staggered, and will depend on the type of entity and the nature and level of risk of the AI system:

- **February 2, 2025**: Prohibited AI systems are banned as of this date. Additionally, providers and deployers of AI systems will be required to ensure sufficient AI literacy for their staff and others dealing with the operation and use of AI systems on this date.

- **May 2, 2025**: Deadline for the EU AI Office to publish the codes of practice to assist providers in demonstrating compliance ahead of their respective deadlines.
- **August 2, 2025**: Entities subject to the Act will be required to comply with obligations applicable to GPAI. Certain enforcement-related provisions will also come into effect, such as penalties for non-compliance.
- **August 2, 2026**: The majority of rules in the EU AI Act will become applicable, including most obligations on high risk and limited risk AI systems.
- **August 2, 2027**: The remaining obligations relating to high risk AI systems that are subject to specified EU product safety legislation listed in Annex I of the EU AI Act, such as regulations governing machinery, toys, recreational crafts, and motor vehicles, will become applicable.

What Are the Potential Enforcement and Penalties Under the Act?

The Act imposes different penalty amounts depending on the provisions that have been violated. For example, entities that fail to comply with the Act's ban on prohibited AI systems may face penalties of the greater of **€35 million or 7% of global annual revenue**. Violations of other provisions (such as those governing provision of high risk AI systems) may result in a maximum fine of the greater of **€15 million or 3% of global annual revenue**. Entities that provide incorrect, incomplete, or misleading information to regulatory authorities may be subject to a fine of the greater of **€7.5 million or 1% of global annual revenue**. These financial penalties are higher than those imposed under the GDPR, and in the case of violations of the EU AI Act that also constitute violations of the GDPR, there is a potential for stacked fines.

What's Next?

EU member states must **designate AI regulators** by August 2, 2025, which may be new entities or existing authorities, and the responsibility may be divided among different regulatory bodies. However, the European Commission's new **AI Office** will have exclusive enforcement powers for GPAI. Three advisory bodies will assist in implementing the EU AI Act:

- The **European Artificial Intelligence Board** will oversee the consistent application of the EU AI Act across EU member states and serve as the primary platform for cooperation between the European Commission and member states.

- A **panel of independent scientific experts** will provide technical advice and input on enforcement, including issuing alerts to the AI Office regarding risks linked to GPAI models.
- Additionally, the **AI Office** will publish the AI Code of Practice detailing AI Act rules for providers of GPAI models, including transparency and copyright disclosure requirements as well as systemic risk taxonomy, assessment and mitigation measures. To develop the Code of Practice, the AI Office will consult an advisory forum made up of industry stakeholders. Once published, the EU Commission may grant the Code of Practice general validity within the EU through an implementing act.

Will the UK Implement a Version of the Act?

The United Kingdom, no longer a member of the EU, will not implement the EU AI Act, and will instead take a different approach to regulation of AI in future legislation. The UK is expected to introduce AI legislation in the middle ground between the executive order-based strategy of the U.S. and the comprehensive approach of the EU AI Act by establishing the **AI Safety Institute**, an independent body, to make voluntary agreements with AI companies. The UK is also expected to amend the UK GDPR and Data Protection Act with additional provisions that may affect companies using or providing AI services. For example, on October 23, 2024, the UK government introduced a draft **Data (Use and Access) Bill** to bolster provisions relating to data subject rights, data processing purpose limitations, international data transfers, automated decision-making technologies, digital verification, and trust certification.

However, note that UK entities can still be subject to the EU AI Act under its extraterritorial applicability. Further, companies using AI to process UK consumer data for the purpose of automated decision making and profiling that affects an individual's legal rights must comply with the UK GDPR, including requirements to comply with data subject rights (e.g., obtaining express consent) and implementation of transparency measures.

Will the U.S. Implement a Version of the Act?

In the U.S., there is currently no comprehensive federal legislation on AI. In lieu of a comprehensive regulation like the EU AI Act, use of AI is currently governed by a mix of federal and state regulations, as well as voluntary industry coalitions.

At the federal level, the Biden Administration's **Executive Order on the Safe, Secure, and Trustworthy Development and Use of AI** delegates responsibility for

issuing AI regulations, providing oversight and enforcing guidelines to its federal agencies, including the White House Office of Management and Budget, Federal Trade Commission, Federal Communications Commission, U.S. Patent and Trademark Office, and U.S. Copyright Office. On October 24, 2024, the White House issued a **National Security Memorandum on AI** that directs the U.S. government to implement steps to: (1) ensure the development of safe, secure, and trustworthy AI, (2) harness AI technologies to advance the U.S. national security mission, and (3) advance international consensus and AI governance. However, with a new president-elect, we will have to wait and see whether the Trump Administration will modify, rescind, or replace President Biden's landmark Executive Order.

At the state level, AI companies must comply with a patchwork of state data privacy and AI regulations. In addition to compliance with state consumer privacy laws affecting certain AI-powered data processing activities, companies must also comply with new AI-centric regulations addressing harms such as algorithmic discrimination, AI watermarking, training data transparency, election fraud, and infringement of IP and privacy rights.

Please refer to our latest publication on the current U.S. AI regulatory landscape here.

How Can GD Help?

Gunderson Dettmer is committed to fostering AI education for the innovation economy and we will continue to monitor and report on issues relating to the EU AI Act that may impact your business, so please stay tuned for further updates. In the meantime, please refer to the Gunderson Dettmer Generative AI Resources page for additional educational materials and insights.

If you have any questions regarding this client alert, or if your company needs assistance evaluating its obligations under the EU AI Act, please reach out to your Gunderson Dettmer attorney.

Related People

Katherine S. Gardner
PARTNER
P +1 212 430 3188

Aaron G. Rubin
PARTNER
P +1 212 430 3181

Anna C. Westfelt
PARTNER
P +1 650 463 5367

Related Services

AI & Machine Learning

Data Privacy

Strategic Transactions & Licensing

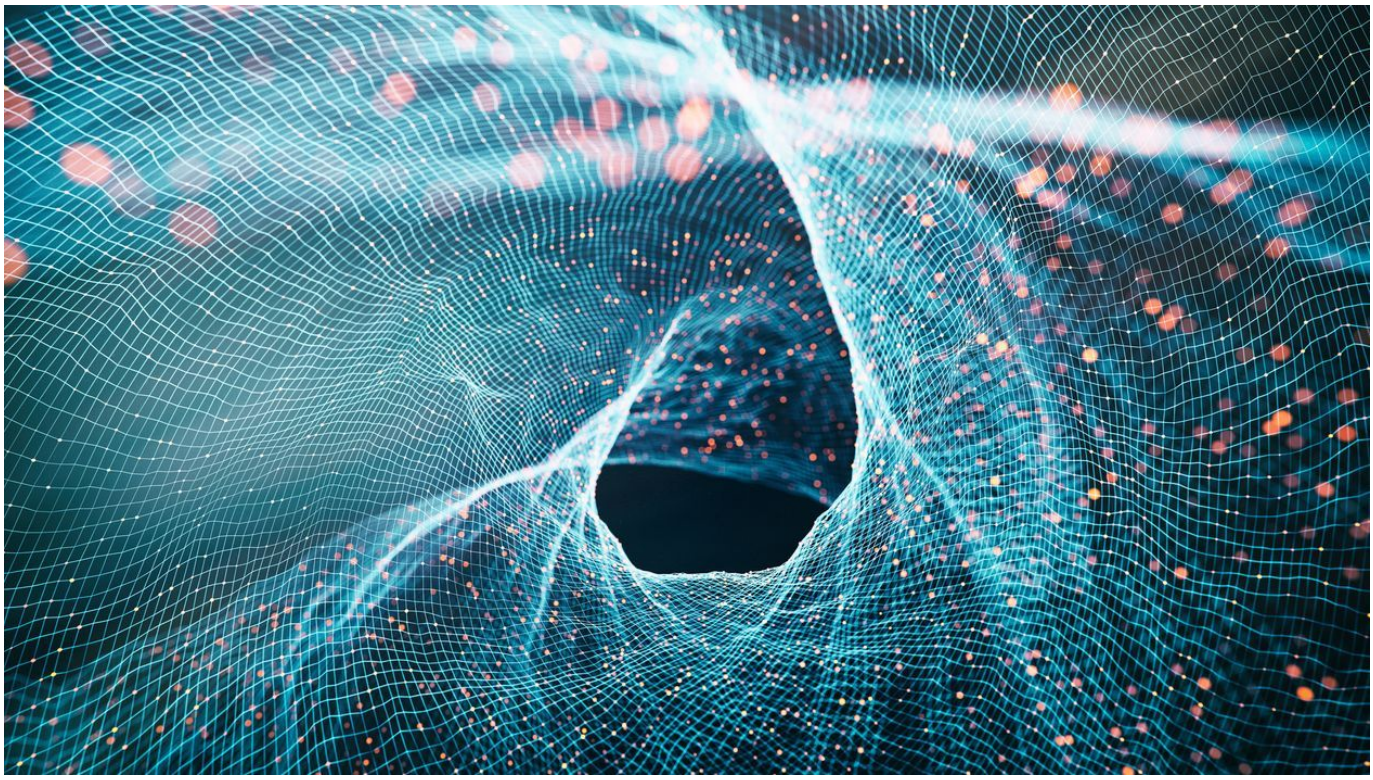

AI @ GD

Gunderson Dettmer’s Generative AI Resources

Gunderson Dettmer is committed to fostering AI education for the innovation economy by supporting startups and venture capital firms.

Discover our AI-focused resources designed to provide updates, education, and insights into the development of AI and generative AI.

Featured Insights

FIRM NEWS

Gunderson Dettmer Commemorates 2025 Asian American and Pacific Islander Heritage (AAPI) Month

CLIENT NEWS

Brazilian Carbon Capture Company Mombak Announces $30M Financing

CLIENT NEWS

Africa B2B OmniRetail Announces $20M Financing

CLIENT NEWS

Glacier Announces Series A Financing to Expand Robot Recycling Fleet

CLIENT NEWS

Dataminr Announces $100M Investment Led by Fortress Investment Group

CLIENT NEWS

Omnidian Announces $87M Series C for Renewable Energy Performance

INSIGHTS

Splitting the Pie: How Savvy Founders Divide Ownership and Navigate Other Founder Equity Decisions

CLIENT NEWS

Chainguard Announces $356 Million Series D Led by Kleiner Perkins and IVP

INSIGHTS

Client Insight: California AI Transparency Act

INSIGHTS

Client Insight: Prepare for BE-10 Benchmark Survey of US Direct Investment Abroad

INSIGHTS

Tech Brew Interviews Aaron Rubin in “Where the legal battle stands around copyright and AI training”

CLIENT NEWS

Latin America Fintech Belvo Announces $15M Funding