



In what could be a harbinger of the future regulation of artificial intelligence (AI) in the United States, the European Commission published its recent [proposal for regulation of AI systems](#). The proposal is part of the European Commission's larger [European strategy for data](#), which seeks to "defend and promote European values and rights in how we design, make and deploy technology in the economy." To this end, the proposed regulation attempts to address the potential risks that AI systems pose to the health, safety, and fundamental rights of Europeans caused by AI systems.

Under the proposed regulation, AI systems presenting the least risk would be subject to minimal disclosure requirements, while at the other end of the spectrum AI systems that exploit human vulnerabilities and government-administered biometric surveillance systems are prohibited outright except under certain circumstances. In the middle, "high-risk" AI systems would be subject to detailed compliance reviews. In many

cases, such high-risk AI system reviews will be in addition to regulatory reviews that apply under existing EU product regulations (e.g., the EU already requires reviews of the safety and marketing of [toys](#) and [radio frequency devices](#) such as smart phones, Internet of Things devices, and radios).

Applicability

The proposed AI regulation applies to all providers that market in the EU or put AI systems into service in the EU as well as users of AI systems in the EU. This scope includes governmental authorities located in the EU. The proposed regulation also applies to providers and users of AI systems whose output is used within the EU, even if the producer or user is located outside of the EU. If the proposed AI regulation becomes law, the enterprises that would be most significantly affected by the regulation are those that provide high-risk AI systems not currently subject to detailed compliance reviews under existing EU product regulations, but that would be under the AI regulation.

Scope of AI Covered by the AI Regulation

The term "AI system" is defined broadly as software that uses any of several identified approaches to generate outputs for a set of human-defined objectives. These approaches cover far more than artificial neural networks and other technologies currently viewed by many as traditional as "AI." In fact, the identified approaches cover many types of software that few would likely consider "AI," such as "statistical approaches" and "search and optimization methods." Under this definition, the AI regulation would seemingly cover the day-to-day tools of nearly every e-commerce platform, social media platform, advertiser, and other business that rely on such commonplace tools to operate.

This apparent breadth can be assessed in two ways. First, this definition may be intended as a placeholder that will be further refined after the public release. There is undoubtedly no perfect definition for "AI system," and by releasing the AI regulation in its current form, lawmakers and interested parties can alter the scope of the definition following public commentary and additional analysis. Second, most "AI systems" inadvertently caught in the net of this broad definition would likely not fall into the high-risk category of AI systems. In other words, these systems generally do not negatively affect the health and safety or fundamental rights of Europeans, and would only be subject to disclosure obligations similar to the data privacy regulations already applicable to most such systems.

Prohibited AI Systems

The proposed regulation prohibits uses of AI systems for purposes that the EU considers to be unjustifiably harmful. Several categories are directed at private sector actors, including prohibitions on the use of so-called "dark patterns" through "subliminal techniques beyond a person's consciousness," or the exploitation of age,

physical or mental vulnerabilities to manipulate behavior that causes physical or psychological harm.

The remaining two areas of prohibition are focused primarily on governmental actions. First, the proposed regulation would prohibit use of AI systems by public authorities to develop "social credit" systems for determining a person's trustworthiness. Notably, this prohibition has carveouts, as such systems are only prohibited if they result in a "detrimental or unfavorable treatment," and even then only if unjustified, disproportionate, or disconnected from the content of the data gathered. Second, indiscriminate surveillance practices by law enforcement that use biometric identification are prohibited in public spaces except in certain exigent circumstances, and with appropriate safeguards on use. These restrictions reflect the EU's larger concerns regarding government overreach in the tracking of its citizens. Military uses are outside the scope of the AI regulation, so this prohibition is essentially limited to law enforcement and civilian government actors.

High-Risk AI Systems

"High-risk" AI systems receive the most attention in the AI regulation. These are systems that, according to the memorandum accompanying the regulation, pose a significant risk to the health and safety or fundamental rights of persons. This boils down to AI systems that (1) are a regulated product or are used as a safety component for a regulated product like toys, radio equipment, machinery, elevators, automobiles, and aviation, or (2) fall into one of several categories: biometric identification, management of critical infrastructure, education and training, human resources and access to employment, law enforcement, administration of justice and democratic processes, migration and border control management, and systems for determining access to public benefits. The regulation contemplates this latter category evolving over time to include other products and services, some of which may face little product regulation at present. Enterprises that provide these products may be venturing into an unfamiliar and evolving regulatory space.

High-risk AI systems would be subject to extensive requirements, necessitating companies to develop new compliance and monitoring procedures, as well as make changes to products both on the front end and the back end such as:

- Developing and maintaining a risk management system for the AI system that considers and tests for foreseeable risks in the AI system;
- Creating extensive technical documentation, such as software architecture diagrams, data requirements, descriptions of how the enterprise built and selected the model used by the AI system, and the results of examinations of the training data for biases;
- Ensuring the security of data and logging of human interactions for auditing purposes;
- Providing detailed instructions to the user about the provider of the AI system, the foreseeable risks with use of the AI system, and the level of accuracy and robustness of the AI system;
- Ensuring the AI system is subject to human oversight (which can be delegated to the user), including a "stop" button or similar procedure;
- Undergoing pre-release compliance review (internal or external based on category of AI system), and post-release audits; and
- Registering the system in a publicly accessible database.

Transparency Requirements

The regulation would impose transparency and disclosure requirements for certain AI systems regardless of risk. Any AI system that interacts with humans must include disclosures to the user they are interacting with an AI system. The AI regulation provides no further details on this requirement, so a simple notice that an AI system is being used would presumably satisfy this regulation. Most "AI systems" (as defined in the regulation) would fall outside of the prohibited and high-risk categories, and so would only be subject to this disclosure obligation. For that reason, while the broad definition of "AI system" captures much more than traditional artificial intelligence techniques, most enterprises will feel minimal impact from being subject to these regulations.

Penalties

The proposed regulation provides for tiered penalties depending on the nature of the violation. Prohibited uses of AI systems (subliminal manipulation, exploitation of vulnerabilities, and development of social credit systems) and prohibited development, testing, and data use practices could result in fines of the higher of either 30,000,000 EUR or 6% of a company's total worldwide annual revenue. Violation of any other requirements or obligations of the proposed regulation could result in fines of the higher of either 20,000,000 EUR or 4% of a company's total worldwide annual revenue. Supplying incorrect, incomplete, or misleading information to certification bodies or national authorities could result in fines of the higher of either 10,000,000 EUR or 2% of a company's total worldwide annual revenue.

Notably, EU government institutions are also subject to fines, with penalties up to 500,000 EUR for engaging in prohibited practices that would result in the highest fines had the violation been committed by a private actor, and fines for all other violations up to 250,000 EUR.

Prospects for Becoming Law

The proposed regulation remains subject to amendment and approval by the European Parliament and potentially the European Council, a process which can take several years. During this long legislative journey, components of the regulation could change significantly, and it may not even become law.

Key Takeaways for U.S. Companies Developing AI Systems

Compliance With Current Laws

Although the proposed AI regulation would mark the most comprehensive regulation of AI to date, stakeholders should be mindful that current U.S. and EU laws already govern some of the conduct it attributes to AI systems.

For example, U.S. federal law prohibits unlawful discrimination on the basis of a protected class in numerous scenarios, such as in employment, the provision of public accommodations, and [medical treatment](#). Uses of AI systems that result in unlawful discrimination in these arenas already pose significant legal risk. Similarly, AI systems that affect public safety or are [used in an unfair or deceptive manner](#) could be regulated through existing consumer protection laws.

Apart from such generally applicable laws, U.S. laws regulating AI are limited in scope, and focus on disclosures related to [AI systems interacting with people](#) or are limited to providing guidance under current law in an industry-specific manner, such as with [autonomous vehicles](#). There is also a movement towards enhanced transparency and disclosure obligations for users when their personal data is processed by AI systems, as discussed further below.

Implications for Laws in the United States

To date, no state or federal laws specifically targeting AI systems have been successfully enacted into law. If the proposed EU AI regulation becomes law, it will undoubtedly influence the development of AI laws in Congress and state legislatures, and potentially globally. This is a trend we saw with the EU's General Data Protection Regulation (GDPR), which has shaped new data privacy laws in California, Virginia, Washington, and several bills before Congress, as well as laws in other countries.

U.S. legislators have so far proposed bills that would regulate AI systems in a specific manner, rather than comprehensively as the EU AI regulation purports to do. In the United States, ["algorithmic accountability" legislation](#) attempts to address concerns about high-risk AI systems similar to those articulated in the EU through self-administered impact assessments and required disclosures, but lacks the EU proposal's outright prohibition on certain uses of AI systems, and nuanced analysis of AI systems used by government actors. Other bills would solely regulate government procurement and use of AI systems, for example, [California AB-13](#) and [Washington SB-5116](#), leaving industry free to develop AI systems for private, nongovernmental use. Upcoming privacy laws such as the [California Privacy Rights Act](#) (CPRA) and the [Virginia Consumer Data Protection Act](#) (CDPA), both effective January 1, 2023, do not attempt to comprehensively regulate AI, instead focusing on disclosure requirements and data subject rights related to profiling and automated decision-making.

Conclusion

Ultimately, the AI regulation (in its current form) will have minimal impact on many enterprises unless they are developing systems in the "high-risk" category that are not currently regulated products. But some stakeholders may be surprised, and unsatisfied with, the fact that the draft legislation puts relatively few additional restrictions on purely private sector AI systems that are not already subject to regulation. The drafters presumably did so to not overly burden private sector activities. But it is yet to be seen whether any enacted form of the AI regulation would strike that balance in the same way.

© 2021 Perkins Coie LLP

Authors



Marc S. Martin

Partner

MMartin@perkinscoie.com [202.654.6351](tel:202.654.6351)



Kyle R. Canavera

Partner

KCanavera@perkinscoie.com [858.720.5782](tel:858.720.5782)

Explore more in

[Consumer Protection](#) [Technology Transactions & Privacy Law](#) [Privacy & Security](#) [Artificial Intelligence & Machine Learning](#) [Communications](#) [Blockchain, Digital Assets & Custody](#)

Related insights

Update

[Coming Soon: Judicial and Agency Interpretations of Washington's Pay Disclosure Law](#)

Update

[October Tip of the Month: U.S. Department of Labor Issues Guidance on AI in the Workplace](#)