AI Legislative and Regulatory Efforts Pick Up Steam: What We're Watching

Alysa Z. Hutnik, Ioana Gorecki, Alexander I. Schneider

July 3, 2024

AI capabilities are growing by the day, and with them, so are increasing government efforts to put in place guardrails, principles, and rules to govern the AI space. In May alone, Utah's Artificial Intelligence Policy Act became the first state-level AI law to take effect, Colorado and Minnesota enacted new laws addressing AI, and the European Union passed historic comprehensive AI regulations. Meanwhile, the FTC continues to issue AI-related guidance materials that emphasize the importance of transparency in human-AI interactions, especially those involving native advertising (prior guidance here and here). As we continue to monitor the flurry of activity underway, we outline below new laws and important bills, standards, and initiatives to monitor.

Federal Efforts

American Privacy Rights Act

Last week, the House Energy and Commerce Committee abruptly canceled a scheduled markup of the latest American Privacy Rights Act (APRA) discussion draft, Congress's most recent comprehensive privacy proposal. Some privacy advocates welcomed the cancellation, strongly opposing the removal of AI and civil rights protections in the latest draft. These protections included prohibitions against algorithmic discrimination and requirements for transparency and impact assessments for AI systems.

At present, it seems APRA may not advance as far as the 2022 American Data Privacy and Protection Act, which was passed out of the Energy and Commerce Committee but ultimately never received a floor vote. With the August recess and October break ahead of the November elections approaching, the likelihood of any comprehensive privacy legislation reaching the House floor this year seems dim. However, we will continue to monitor these federal legislative efforts and their potential impact on AI providers.

White House Executive Order

Last year, the White House released the federal government's first comprehensive guidelines regarding AI. Although the Executive Order focuses almost entirely on the government's own use of AI, the ultimate effects of the order will be significant for private sector businesses engaging with federal agencies.

Pursuant to the Executive Order, on April 29, 2024, NIST released a draft risk management profile specifically addressing generative AI. The Generative AI Profile—which is intended as a companion resource to NIST's AI Risk Management Framework—offers voluntary best practice guidance

regarding the design, deployment, and operation of generative AI systems. As states continue to draft AI legislation, the NIST AI Risk Management Framework will likely continue to serve as an instructive reference point for legislators across the country.

State Legislation

Colorado AI Act

The Colorado AI Act, SB 205, is now set to take effect February 1, 2026, although the freshly-signed law is already slated for revisions: in a recent letter, Gov. Jared Polis, AG Phil Weiser and Senate Majority Leader Robert Rodriguez acknowledged that "a state by state patchwork of regulation" on AI poses "challenges to the cultivation of a strong technology sector" and promised to engage in a process to revise the new law to "minimize unintended consequences associated with its implementation."

As drafted, the law introduces new obligations and reporting requirements for both developers and deployers of AI systems. Key requirements include:

- **Transparency.** Moving forward, any businesses that use AI systems to interact with consumers must disclose this fact during consumer interactions.
- **Algorithmic Discrimination in High-Risk AI Systems.** The new law seeks to combat "algorithmic discrimination," where the use of AI results in outcomes that disfavor consumers based on several personal and sensitive data categories. High-risk AI systems are defined as systems used to make decisions about individuals in the areas of education, employment, finance or lending, government services, healthcare, housing, insurance, and legal. Developers and deployers of such systems have a duty to use reasonable care to protect consumers from known or reasonably foreseeable risks of algorithmic discrimination, and the law identifies specific obligations such entities must undertake.
- **Consumer Notice, Correction, and Opt-Out Rights**. Consumers must be notified when high-risk AI systems are used to make any decisions about them in the areas outlined above (e.g., education, employment, etc.), and must have the right to correct inaccurate data and appeal the decision to a human reviewer.
- **Existing Obligations Under the Colorado Privacy Act (CPA).** Deployers must also respect the existing rights of consumers under the CPA, including the right to opt-out of the processing of personal information for profiling with legal or similarly significant effects concerning the consumer, including decisions made using AI. In April, Colorado amended the CPA's definition of sensitive data to include both biological and neural data used either in insolation or in combination with other personal data elements for identification purposes. The CPA additionally creates AI-related disclosure obligations, requiring businesses to provide privacy policy language that details the personal data categories used for profiling, a plain-language explanation regarding the AI logic in use, explanations describing its benefits and potential consequences, and text explaining whether the system has been evaluated for accuracy, fairness or bias.
- **Enforcement**. The Colorado attorney general has sole authority to enforce the Colorado AI Act, and the law includes no private right of action. Violations are considered breaches of Colorado's general consumer protection laws, which can result in a maximum civil penalty of $20,000 per violation. Notably, each violation is counted individually for every affected consumer or

transaction. Consequently, just 50 impacted consumers could result in a maximum civil penalty of $1 million. Actions must be brought within three years of the violation occurring, or from the time when the violation was discovered.

We'll keep an eye on whether all these requirements survive the revision process suggested above.

Utah Artificial Intelligence Policy Act

On May 1, 2024, Utah's Artificial Intelligence Policy Act, SB 149, became effective. Generally, Utah's legislature has pursued a far lighter touch to AI regulation than Colorado. Key takeaways include:

- **Disclosure Upon Request.** Most businesses and individuals will only be required to disclose the use of AI when prompted by a consumer.
- **Disclosing the Use of AI in Regulated Professions.** Businesses and individuals operating within regulated professions (e.g., healthcare professionals) must prominently disclose the use of AI before its use with customers.
- **Responsibility for Generative AI Outputs**. Companies are responsible for the outputs of their generative AI tools and cannot pass on blame if those tools violate Utah consumer protection laws.

Comprehensive State Privacy Laws

Twenty states have now passed comprehensive state privacy laws: California, Colorado, Connecticut, Delaware, Florida, Indiana, Iowa, Kentucky, Maryland, Minnesota, Nebraska, New Hampshire, New Jersey, Oregon, Rhode Island, Tennessee, Texas, Utah, and Virginia. These states, with the exceptions of Utah and Iowa, impose additional requirements on companies engaging in "profiling," which is defined as the automated processing of personal data to analyze or predict something personal about an individual, such as one's economic situation, behavior, health, or personal preferences. Under these laws, consumers must be able to opt-out of being profiled in a manner that could lead to a "legal effect" on that consumer or another "similarly significant effect." Although a few of these laws are currently effective, the majority come into effect over the next few years. Here are the key dates to keep mind:

- **Currently Effective.** The comprehensive privacy laws in California, Colorado, Connecticut, Utah, and Virginia are currently effective.
- **Effective in 2024.** Florida, Montana, Oregon, and Texas have comprehensive privacy laws coming into effect in the next several months.
- **Effective in 2025.** Delaware, Iowa, Maryland, Minnesota, Nebraska, New Hampshire, New Jersey, and Tennessee have enacted comprehensive data privacy laws that will become effective in 2025.
- **Effective in 2026.** Kentucky and Indiana have enacted comprehensive data privacy laws that will become effective on Jan. 1, 2026. The Rhode Island legislature also passed the Rhode Island Data Transparency and Privacy Protection Act, SB 2500 / HB 7787, on June 13, 2024. If signed, the law will also become effective on Jan. 1, 2026.

California Privacy Protection Agency Initiatives

The California Privacy Protection Agency is currently considering rules and engaging in pre-formal

rulemaking stakeholder sessions regarding the use of automated decision making technology (ADMT). California defines ADMT as technology that collects, uses, retains or discloses personal information and either replaces or substantially facilitates human decision making. Algorithmic "profiling," discussed above, is encompassed within this definition. Examples include resume-screening tools used by businesses to decide whether to interview applicants and analytics tools that place consumers into audience groups to further target them with advertising.

Businesses subject to the California Consumer Privacy Act (CCPA) and that use ADMT for "extensive profiling," to make "significant decisions" regarding consumers, or that use personal information to train ADMT would be subject to new transparency and opt-out requirements. Behavioral advertising, the practice of tracking users' online activities to deliver ads tailored to their interests, is included within the definition of "extensive profiling." Further discussion regarding the terms "extensive profiling" and "significant decisions" can be found here. Businesses would be required to offer a pre-use notice informing consumers of how the company uses ADMT and of the individual's CCPA opt-out rights.

Ongoing Legislative Efforts

Currently, a multitude of states, including New York, California, and Massachusetts, are working on proposed AI governance bills. In addition, new legislation in Illinois addressing AI usage currently awaits the Governor's signature.

- **California**. The Assembly recently advanced multiple bills addressing AI usage. These bills include provisions prohibiting algorithmic discrimination and would establish new compliance and reporting requirements for AI providers. Additionally, these bills would require businesses to implement watermarking systems identifying AI-generated content and to publicize information regarding the methods used to train AI models.

- **Illinois.** On May 24, 2024, the Illinois legislature passed HB 3773, amending the Illinois Human Rights Act by adding new provisions regarding the use of predictive data analytics for employment and credit decisions.

Europe

The EU AI Act

On May 21, 2024, the EU Council unanimously passed the EU AI Act (AIA). Businesses, whether EU-based or not, should pay close attention to the upcoming changes for two reasons. First, the AIA applies to all providers of AI systems placed on the EU market, regardless of where the provider is located. Second, the penalties for non-compliance are some of the toughest in the world, allowing for fines up to €35 million EUR or 7% of a company's annual revenue.

Broadly, the AIA creates a risk classification scheme, which places AI systems into one of several categories. The categories are:

- **Unacceptable Risk.** AI systems constituting an unacceptable risk are prohibited entirely. These include systems used to manipulate or exploit individuals, classify or evaluate individuals based upon their personal traits, and emotion-recognition systems used in workplace and educational contexts.

- **High Risk.** The AIA defines high risk systems as those presenting a significant risk to health, safety, or fundamental rights. Examples of AI systems falling under this category include those

used in education, employment, healthcare, and banking settings. Providers of high-risk systems are subject to a number of strict regulations, including required registration in a public EU database. Additionally, providers of these systems must perform regular impact assessments and implement procedures that ensure transparency, security, and human oversight of their systems.

- **Limited Risk.** For systems posing limited risks, such as chatbots interacting with humans and AI-generated content, the AIA imposes transparency obligations to ensure humans are informed that an AI system was involved. Providers of AI-generated content must ensure it is identifiable as such.
- **Minimal or No Risk.** Minimal-risk AI uses, which present little to no risk to the rights or safety of individuals, can be freely used under the AIA. Examples include AI-enabled video games and spam filters. Most AI systems currently deployed are likely to fall under this category.
- **General Purpose AI (GPAI)**. GPAI refers to AI systems trained on broad datasets capable of serving a variety of purposes. Popular examples include OpenAI's ChatGPT and DALL-E programs. Providers of GPAI models are required to produce technical documentation and release detailed summaries of their training data. For GPAI models that present systemic risks, providers must also implement cybersecurity measures, mitigate potential risks, and perform evaluations that include adversarial testing.

We will continue to monitor these ongoing state, federal, and international AI legislative efforts and provide you with the latest updates to help you prepare for what lies ahead.

Summer Associate Joe Cahill contributed to this post