

JUNE 15, 2026

# Congress and State Lawmakers Are Racing to Keep Up With AI

---

Goodwin lawyers break down the growing patchwork of state AI laws.

*Gabe Maldoff, Omer Tene, Bethany P. Withers*

States are taking the lead in regulating AI, even as Congress begins to debate a bipartisan frontier AI law and the Trump administration pursues a deregulatory agenda seeking to curb state AI lawmaking through an executive order (EO).

This year is set to break previous records, with more than 1,500 AI bills under consideration in statehouses, many of which are nearing or have crossed the finish line. These laws add to a chorus of state lawmaking, which included 150 AI-related bills that passed into law in 2025, according to a Stanford University report.

Congress's latest foray into AI regulation would, if passed, build on the framework emerging at the state level. Released on June 4, 2026, by Reps. Jay Obernolte (R-California) and Lori Trahan (D-Massachusetts), the Great American AI Act of 2026 would nationalize the frontier-model governance approach emerging in several states. While the bill would preempt certain state laws that regulate AI development for a three-year period, the current discussion draft of the bill would leave much of the current state patchwork intact.

Meanwhile, President Trump's EO, which aims to block state AI lawmaking in order to promote AI innovation and US competitiveness, has not yet displaced the framework emerging at the state level. In fact, a more recent EO, "Promoting Advanced Artificial Intelligence Innovation and Security," adopts a voluntary standard for frontier AI safety similar to the state model and congressional bill.

The structure of AI regulation emerging across these state and federal initiatives does not replicate the European Union's more "comprehensive" AI Act, which prohibits certain uses of AI deemed unacceptable and includes frameworks of varying complexity for "high-risk" and "general purpose" AI systems and models.

Instead, US regulation has focused on several key issues, such as AI companions, mental health applications, the use of AI to make "consequential decisions," generative AI development, and frontier models, as well as AI deployment in critical sectors such as life sciences, healthcare, and financial services.

Here are the key themes underlying the emerging consensus on US AI regulation, as it stands in 2026:

## AI Disclosures for Chatbots

The earliest AI-specific laws — predating the rise of generative AI — required companies to disclose to consumers that they were interacting with AI in limited contexts, such as online sales and political advocacy. As generative AI use cases have proliferated, so have AI transparency measures. Utah’s AI Policy Act, for example, requires those in “regulated occupations” (including healthcare professionals) to disclose to consumers at the beginning of an interaction that they are interacting with generative AI. Outside of regulated occupations, other businesses need to “clearly and conspicuously” tell consumers that they are interacting with generative AI systems when asked.

At their core, these laws are designed to ensure that consumers know when they are interacting with AI rather than a human. Expect that requirements to disclose AI use will continue to proliferate.

## Self-Harm and Children’s Safety Guardrails for Companion AI

Some states, including California, Connecticut, New York, Oregon, and Washington, have begun regulating “companion AI” services that simulate ongoing human-like relationships. These laws reach beyond chatbots that are explicitly marketed as AI companions and could include a wide range of chatbots that retain and draw from conversation logs across user sessions.

Companion AI laws are designed to address emotional manipulation, mental health risks, and vulnerable populations, such as minors, using the services.

Companion AI laws generally require companies to inform (and, in some instances, remind) users that they are engaging with AI and to implement policies and protections against self-harm and suicidal ideation. California’s companion AI law, for instance, mandates that companion chatbot operators publish a safety protocol on their website and refer users to crisis service providers, such as a suicide hotline or crisis text line, if the user expresses signs of suicidal ideation, suicide, or self-harm.

Chatbots offered to minors are subject to more stringent protections. For example, Oregon SB 1546 requires that covered chatbot operators employ reasonable measures to prevent the AI companion from producing sexually explicit content and remind minor users to take a break at least every three hours.

## Restrictions on Mental Health and Therapeutic Uses

States are beginning to draw clearer lines around the use of AI in therapy, with laws including Illinois’s HB 1806 setting an early benchmark.

Most significantly, HB 1806 prohibits AI from replacing licensed therapists in delivering care — complicating the use of more autonomous, agentic AI systems capable of making therapeutic decisions and raising critical questions for use of AI in health and wellness applications, which are among the most popular consumer uses of AI.

Illinois’s law also requires licensed providers to obtain patient consent for certain “supplementary support” uses, including recording or transcribing sessions and maintaining patient file notes, and allows AI to support administrative tasks, such as scheduling therapy sessions, without patient consent.

## Transparency, Explainability, and Rights to Contest Use of AI for Consequential Decisions in Employment, Financial Services, Healthcare, Education, and Housing

Although regulation of automated decision-making technology (ADMT) existed long before the advent of generative AI — including in EU data protection laws dating back to 1995 — the growing sophistication and scale of modern AI systems has brought renewed regulatory focus to these tools used in “consequential” or “significant” decisions affecting employment, financial services, healthcare, education, and housing.

States such as Colorado, Texas, and California have advanced laws and regulations that impose obligations on companies using ADMT. Illinois and New York City have passed similar measures focused specifically on employment. Importantly, these laws apply not just to “solely” automated decisions but also to AI-assisted decisions in which ADMT plays a role in the outcome, though the threshold for triggering regulation varies, from a “substantial factor” standard to the lower “non-de minimis factor” test adopted by Colorado’s amended law.

Common requirements in these laws include notifying consumers when AI is being used in consequential decision-making, giving individuals opportunities to opt out or contest the decision, allowing individuals to seek human review, and requiring companies to adopt risk management policies to minimize error, bias, and discrimination.

Colorado’s recent legislative pivot illustrates how ADMT regulation is evolving. In 2026, the state replaced its landmark Colorado AI Act with the Colorado ADMT Act, broadening the law’s focus from “artificial intelligence” to “automated decision-making technology.” The new law expands coverage to a wider range of digital technologies, and while the amendment replaces many of the original law’s governance requirements with a framework centered on transparency, consumer rights, and accountability, a new liability regime is designed to encourage companies to implement internal procedures and safeguards — similar to those mandated in the original law — to mitigate legal risks.

## Training Data Transparency and Watermarking for Generative AI Developers

Some laws specifically regulate generative AI systems that generate new content based on patterns learned from large training datasets.

California AB 2013, for example, requires developers of generative AI systems, as well as entities that “substantially modify” such systems, to publicly post documentation about the data used to train their models, including the sources and ownership of the datasets, the purpose and methodology of data collection, the inclusion of personal information, and the licensing status of such datasets.

Under another California law, SB 942, developers of generative AI systems with more than 1 million monthly users will need to watermark AI-generated content to facilitate detection and provide tools that allow users to verify whether content was generated by the developer’s system.

## Frontier AI Safety and Safeguards Against Catastrophic Risks

Some states are beginning to address the risks associated with the most advanced AI systems, often

referred to as “frontier models” — defined as foundation models trained on broad datasets, adapted to a wide range of tasks, and using a quantity of computing power greater than  $10^{26}$  computational operations. Both California and New York have adopted laws that require developers of frontier models to implement frameworks designed to protect against “catastrophic risks.”

Frontier AI safety frameworks must include alignment with national, international, or industry standards; internal governance programs; third-party risk assessments; and transparency reporting of such frameworks and assessments. Frontier developers also need to disclose to regulators certain safety incidents within set timelines. A bill in Illinois (SB 315), which has passed the legislature and is now awaiting the governor’s signature, would build on the New York and California laws by also requiring independent third-party auditing of compliance with internal safety frameworks.

The White House is also stepping in to formally oversee advanced AI developers. A recent Trump administration EO asks AI companies to share models with the government for cybersecurity testing before public release. The Great American AI Act of 2026, currently being debated in Congress, would impose transparency reports, a published frontier AI framework, critical safety incident reporting, and whistleblower protections on large-scale frontier developers — measures that largely align with the California and New York laws. The proposed federal law would also incorporate independent auditing requirements similar to Illinois SB 315.

## Sector-Specific Legislation: Employment, Healthcare, Financial and More

Rather than enacting a single omnibus statute, many states are increasingly regulating AI sector by sector.

- In employment, legislative measures target automated employment decision tools (AEDTs) and worker monitoring. Starting with New York City’s Local Law 144 of 2021, which requires bias audits, AEDT laws proliferate, with California beginning to regulate ADMT in employment under the CCPA and Illinois’ amended Human Rights Act now restricting AI-driven hiring decisions.
- In healthcare, states are imposing disclosure and human-oversight requirements: California’s AB 489 prohibits AI from falsely claiming healthcare licenses and requires disclosures when AI communicates with patients, Texas SB 1188 governs use of AI by healthcare practitioners, Nevada’s AB 406 regulates the use of AI for mental and behavioral healthcare, Illinois’ HB 1806 bars chatbots from representing themselves as licensed professionals, and companion chatbot laws like California’s SB 243 mandate suicide- and self-harm-detection safeguards.
- In finance, laws govern the deployment of AI tools to evaluate consumer credit and lending, such as New York’s A773-B, which awaits the governor’s signature.

## Where State AI Laws Are Headed Next

States are moving quickly, and not every use of AI will fit neatly into the above categories. But one theme is clear: Lawmakers are trying to identify the areas where AI can cause real-world harm and put guardrails around them.

For companies, the big task is recognizing the patterns emerging across states and thinking about how their AI systems are developed, deployed, and monitored before regulators force the issue. That challenge is

likely to grow even more complicated as AI systems become increasingly autonomous. Agentic AI systems, which can independently plan and execute tasks, are already raising new questions that many current laws don't fully address.

States are also beginning to scrutinize algorithmic pricing systems that use AI and personal data to automatically adjust prices or offers for consumers. Although algorithmic pricing has existed for years, lawmakers are increasingly concerned that AI dramatically increases the scale, speed, and precision of these systems, creating risks in areas such as antitrust.

As the pace of AI development and adoption accelerates, state AI regulation will remain a rapidly evolving frontier. The question now is not whether AI will be regulated, but rather, whether those developing and deploying AI technologies can keep up.

*This informational piece, which may be considered advertising under the ethical rules of certain jurisdictions, is provided on the understanding that it does not constitute the rendering of legal advice or other professional advice by Goodwin or its lawyers. Prior results do not guarantee similar outcomes.*

## CONTACTS

**Gabe Maldoff**

Partner  
Data, Privacy & Cybersecurity  
gmaldoff@goodwinlaw.com

**Omer Tene**

Partner  
Technology  
otene@goodwinlaw.com

**Bethany P. Withers**

Partner  
Technology  
bwithers@goodwinlaw.com