

Texas enacts TRAIGA with intent-based AI bans and a 36-month sandbox

Source copy kept by wheresthe.ai for the story "Texas enacts TRAIGA with intent-based AI bans and a 36-month sandbox" (<https://wheresthe.ai/d/us-tx-traiga-hb149-signed-2025-06/>).

Source: <https://www.dlapiper.com/insights/publications/2025/06/texas-adopts-the-responsible-ai-governance-act>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 2a99da07f841d992b393da888e4fd38d31f47b92f44348c0cef0ff20380e97d1

United States Home Insights AI Alert: Texas adopts the Responsible AI Governance Act

25 June 2025 • 12 minute read

share

LinkedIn

Facebook

Twitter

Email

Whatsapp

Line

WeChat

Xing

Texas adopts the Responsible AI Governance Act

Written by : Ashley Carr Sean Fulton Karley Buckley Kyle Kloeppel

On June 22, 2025, Governor Abbott signed HB 149, the "Texas Responsible Artificial Intelligence Governance Act" (the Act or TRAIGA), making Texas the third US state, after Colorado and Utah, to adopt a comprehensive artificial intelligence (AI) law. While the law has some similarities to other state and global AI regulations, it also has key differences intended to strike a more innovation-friendly balance, such as a higher "intent" standard for showing algorithmic bias, several safe harbors, and preemption of local AI regulation.

Once the Act takes effect on January 1, 2026, Texas will:

Establish baseline duties for AI "developers" and "deployers,"

Prohibit AI intended for social scoring or discrimination,

Create a first-in-the-nation AI regulatory sandbox,

Vest exclusive enforcement authority with the Attorney General (AG), and

Broadly pre-empt local ordinances governing AI.

The Act's effective date gives companies roughly 18 months to stand up compliance programs. Organizations that operate nationally or globally are encouraged to harmonize those efforts with the fast-approaching requirements of the EU AI Act and Colorado's AI Act. Companies are also encouraged to monitor federal negotiations that would tie federal funding to a states' compliance with a ten-year moratorium on AI regulations, which could upend the patchwork altogether.

Below, we summarize the key provisions, compare the Texas framework with other leading regimes, and outline practical steps for organizations building or deploying AI in – or involving – Texas.

Broad scope, familiar definitions

The Act applies to any person that "promotes, advertises, or conducts business" in Texas, offers products or

services to Texas residents, or “develops or deploys” an AI system in the state. Mirroring the EU AI Act formulation, the Act defines an “artificial intelligence system” as “any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs ... that can influence physical or virtual environments.” Importantly, the definition is technology- and sector-agnostic, ensuring coverage of generative models, recommendation engines, biometric systems, and more.

Similar to the EU AI Act and Colorado AI Act, the Texas Act assigns responsibility by role:

Developer: Any person who creates an AI system offered, sold, leased, or otherwise provided in Texas.

Deployer: Any person who puts an AI system into service or use in the state.

Baseline duties and prohibitions

Transparency to consumers: Governmental entities must disclose to individuals that they are interacting with an AI system before or at the point of interaction. The notice must be clear, conspicuous, and in plain language; hyperlinks are expressly permitted.

Manipulation of human behavior: Developers and deployers may not intentionally use AI to incite self-harm, violence, or criminal activity.

“Social scoring” ban: Echoing Article 5 of the EU AI Act, Texas bars government entities from using or deploying AI systems that categorize individuals to assign a “social score” that could lead to detrimental treatment unrelated or disproportionate to the context or gravity of the behaviors that were observed, or that violate US or Texas constitutional rights.

Biometric data guardrails: Without an individual’s consent, a governmental entity may not use AI to uniquely identify a person via biometric data obtained from publicly available sources if doing so would violate any rights under the US or Texas Constitution, or state or federal law.

Constitutional rights and unlawful discrimination: AI systems may not be developed or deployed “with the sole intent” to infringe, restrict, or otherwise impair an individual’s rights under the US Constitution or with the intent to “unlawfully discriminate against a protected class in violation of state or federal law.” A mere disparate impact is insufficient to establish an intent to discriminate against a protected class.

Sexually explicit content involving minors: AI systems may not be developed or deployed “with the sole intent of producing, assisting or aiding in producing, or distributing” AI-generated child sexual abuse material and “deepfake” sexual content depicting minors.

Notably, the Act sets a high threshold for liability with respect to two of the key prohibitions outlined above. Specifically, to violate the Act’s prohibitions on using AI to infringe on an individual’s rights under the Constitution or to unlawfully discriminate against a protected class under state or federal law, a developer or deployer must have acted with intent. Requiring intent to establish liability for violations involving constitutional or anti-discrimination protections makes the Texas Act unique among other AI-specific state laws.

AI regulatory sandbox

The Act creates a 36-month sandbox administered by the Department of Information Resources (DIR) in consultation with the new Texas Artificial Intelligence Council. Approved participants may test innovative AI applications without obtaining otherwise-required state licenses or permits, subject to:

An application describing the system, benefits, risks, and mitigation measures,

Quarterly reporting of performance metrics and consumer feedback, and

DIR’s power to terminate participation if risks become undue, if the AI violates any federal law or regulation, or if the AI violates any state law or regulation that is not waived under the program.

Crucially, the Texas AG may not “file or pursue charges against a program participant for violation of a law or regulation waived under this chapter that occurs during the testing period.”

Enforcement and pre-emption

Exclusive authority: The Texas AG has exclusive authority to enforce this Act. The AG is required to create and maintain an online mechanism to submit complaints for review.

Notice-and-cure: If the AG determines a person has violated the Act, they must provide notice to cure and cannot sue until 60 days have passed without a cure.

Civil penalties: Penalties include fines of up to \$200,000 per uncurable violation, plus up to \$40,000 per day for continuing violations, with lesser tiers for curable breaches or statements that prove inaccurate.

Statewide pre-emption: The Act expressly nullifies any city or county ordinances regulating AI, aiming to prevent a local patchwork.

No private right of action: The Act does not create a private right of action – again aligning with Colorado and diverging from some state privacy statutes.

Safe harbors

Notably, the Act sets out multiple safe harbors for organizations against which the AG has sought a civil penalty or injunction. For example, a defendant may not be found liable if they “discover a violation ... through testing, including adversarial testing or red-team testing” or if they “substantially compl[y] with the most recent version of the ‘Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile’ published by the National Institute of Standards and Technology or another nationally or internationally recognized risk management framework for AI systems, an internal review process” (emphasis added).

How does Texas compare?

Texas Colorado EU AI Act

Effective date January 1, 2026 February 1, 2026 February 2, 2025 for first obligations; implementation is staggered

Risk framework Duties/prohibitions keyed to specific practices; no formal tiering Framework designed for “High-risk” AI Prohibited, High, Limited, and Minimal

Actors covered Developers and deployers Developers and deployers Providers, deployers, importers, and distributors

Transparency Mandatory AI-interaction notice for government entities Consumer notice for high-risk consequential decisions Mandatory disclosure for most AI interactions

Discrimination standard “Intent to discriminate” required “Algorithmic discrimination” (impact-focused) Fundamental-rights impact assessment for high-risk AI systems

Sandbox Yes – 36 months, broad No Member States must establish at least one AI regulatory sandbox at the national level by August 2, 2026

Penalties AG enforcement; up to USD200,000 AG enforcement; penalties to be determined by rule Tiered penalties by type of violation; up to EUR35 million or 7 percent of global annual turnover

The looming federal moratorium

On May 22, 2025, the US House budget reconciliation bill advanced with a ten-year moratorium that would bar states and localities from enforcing AI-specific laws.

An emerging Senate draft of the bill, released on June 5, 2025 by the Senate Committee on Commerce, Science, and Transportation, would condition receipt of future Broadband, Equity, Access, and Deployment (BEAD) dollars on a state’s agreement not to regulate AI models or automated decision systems. If Congress adopts the funding-lever approach, Texas will confront an immediate policy fork: (i) suspend enforcement of the Act or (ii) proceed with the Act and forfeit federal broadband and infrastructure funds.

Although the Senate version is narrower than the House’s blanket pre-emption, it would still functionally freeze the state’s enforcement mechanisms of this Act.

On June 20, 2025, the provision passed a major hurdle when the Senate Committee on the Budget determined that the provision did not violate the Senate’s Byrd Rule.

Next steps

While uncertainty remains about the impact of state laws in the context of a potential federal moratorium,

companies can take steps to help prepare for enforcement and mature their AI governance strategy, including:

Taking inventory and stratifying AI use cases by risk level,

Documenting cessation of “prohibited” practices under the Act (and similar laws in other jurisdictions),

Establishing testing protocols and template documentation, including for adversarial testing and red teaming, and

Aligning enterprise-wide AI governance with the National Institute of Standards and Technology AI Risk Management Framework and other industry risk management standards, such as ISO 42001.

Find out more

DLA Piper’s team of AI lawyers, data scientists, and policy experts helps organizations navigate the complex workings of their AI systems and comply with current and developing regulatory requirements. We continuously monitor updates and developments arising in AI and its impact on industry across the world.

For more information on AI and the emerging legal and regulatory standards, visit DLA Piper’s focus page on AI .

Gain insights and perspectives that will help shape your AI strategy through our AI ChatRoom series .

For further information or if you have any questions, please contact any of the authors.

Related insights Publication Legal red teaming, one year in: Reports from the field 23 June 2025 . 1 minute read

Publication Landmark AI framework sets new standard for tackling algorithmic bias 6 February 2025 . 4 minute read

Publication Ten-year moratorium on AI regulation proposed in US Congress 22 May 2025 . 10 minute read

Print

Related capabilities Technology Artificial Intelligence and Data Analytics Regulatory and Government Affairs Data, Privacy and Cybersecurity