



ALERT · JULY 9, 2025

With TRAIGA, Lone Star State Leans Into AI Governance Regulation

BY Omer Tene

With the demise of the proposed federal moratorium on state-level AI regulations, which Congress eliminated at the last moment from the “One Big Beautiful Bill” last week, state AI legislation takes center stage. The AI moratorium would have prevented states from instituting AI regulatory measures for a decade. And while businesses might think state AI regulation is the bailiwick of blue states like California, Connecticut, and Colorado — which, to be sure, already have AI laws in the books, if not yet in force — red states too are joining the fray. No red state has more oomph than Texas, and indeed, the Lone Star State has now passed a major AI statute, the Texas Responsible Artificial Intelligence Governance Act (TRAIGA).

TRAIGA (HB 149), the brainchild of Texas Rep. Giovanni Capriglione, a Republican, had been percolating among policymakers for a couple of years. Capriglione, who also authored the state’s robust privacy law, the Texas Data Privacy and Security Act (TDPSA), initially proposed a much stricter measure that resembled a cross between Colorado’s AI Act and, dare we say, the EU AI Act. How times have changed. Immediately after Republicans scored the trifecta, sweeping into control of the White House and both chambers of Congress, in November 2024, TRAIGA went in for repairs. When it emerged in its current form, which Gov. Greg Abbott signed into law on June 22, 2025, it was more accommodating of technology and innovation and less focused on AI risks. Still, TRAIGA cements the protection of civil rights in the age of AI; sets forth significant guardrails for AI systems; introduces important AI policy nomenclature; empowers the Texas attorney general with strong enforcement powers; and innovates with regulatory mechanisms such as an AI sandbox and an AI council. Notably, the law doesn’t provide companies much time to prepare. It kicks in January 1, 2026, even before its older relative from Colorado, which has been in the books for more than a year but doesn’t go into effect until February 1, 2026.

Who It Applies To

Supporters of the federal AI moratorium argued that a federal framework is necessary to regulate a technology that without a doubt affects interstate commerce. Granted, TRAIGA will apply to many businesses that don’t have a foothold in Texas. It will apply to any person who “(1) promotes, advertises, or conducts business in [Texas]; (2) produces a product or service used by residents of this state; or (3) develops or deploys an artificial intelligence system in this state.” This means that a California startup fine-tuning a model later sold to Dallas hospitals will be subject to TRAIGA, as will an Austin developer who sells an AI system to customers in Illinois. Like other AI governance laws, including Colorado’s and the EU’s, TRAIGA distinguishes between a “developer,” which means “a person who develops an artificial intelligence system that is offered, sold, leased, given, or otherwise provided in this state,” and a “deployer,” who is “a person who deploys an artificial intelligence system for use in this state.”

Similar to other AI laws, TRAIGA 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, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments.”

AI Governance

Unlike the EU AI Act, TRAIGA doesn't impose burdensome governance and accountability measures, much less a licensing scheme. However, one of its most important contributions is implicit in one of its enforcement clauses: comprising a list of steps that businesses would be wise to adopt to satisfy emerging governance standards. The law states that when the attorney general's office receives an individual's complaint alleging a TRAIGA violation, it may issue a civil investigative demand (CID). TRAIGA proceeds to specify what type of information the attorney general may seek in a CID, thus providing a compliance road map for businesses to prepare and document on the back end. Surely, you don't want to be caught empty-handed if the Texas attorney general knocks on your door. The list includes:

- A high-level description of the purpose, intended use, deployment context, and associated benefits of the AI system
- A description of the type of data used to program or train the AI system
- A high-level description of the categories of data processed as inputs for the AI system
- A high-level description of the outputs produced by the AI system
- Any metrics the business uses to evaluate the performance of the AI system (mental note: Compile metrics!)
- Any known limitations of the AI system
- A high-level description of the post-deployment monitoring and user safeguards the business uses for the AI system, including, if the person is a deployer, the oversight, use, and learning process established by the person to address issues arising from the system's deployment

Hence, while the law does not impose governance requirements through the front door, it in fact institutes a robust set of recordkeeping and accountability measures, which explicitly apply not just to developers but also to deployers of AI.

Protection of Civil Rights

Make no mistake, TRAIGA is a civil rights law. This is evident from the following:

- TRAIGA's Construction and Application clause specifies, among other purposes, “protect[ing] individuals and groups of individuals from known and reasonably foreseeable risks associated with artificial intelligence systems.”
- The law provides individuals with rights against deployment of AI systems by the state.
- It prohibits government entities from developing a “social scoring” system, meaning an AI system that “evaluates or classifies a natural person or group of natural persons based on social behavior or personal characteristics, whether known, inferred, or predicted, with the intent to calculate or assign a social score or similar categorical estimation or valuation ... that results or may result in” detrimental or unfavorable treatment in a different domain or that is unjustified or disproportionate to the nature or gravity of the observed behavior or characteristics, or the infringement of any federal or state constitutional or legal right.
- It prohibits AI systems from “manipulation of human behavior,” which includes “intentionally aim[ing] to incite or encourage a person to (1) commit physical self-harm, including suicide; (2) harm another person; or (3) engage in criminal activity.”
- It restricts government from rolling out biometric identification systems, except with individual consent.

Most importantly, TRAIGA prohibits the use of AI — in both the public and private sectors — for unlawful discrimination against protected class. The law states, “A person may not develop or deploy an artificial intelligence system with the intent to unlawfully discriminate against a protected class in violation of state or federal law.” It defines a protected class as “a group or class of persons with a characteristic, quality, belief, or status protected from discrimination by state or federal civil rights laws, and

includes race, color, national origin, sex, age, religion, or disability.” That’s a powerful message sent by the Texas legislature to developers and deployers of AI. This clause has certain exemptions for financial institutions, including insurance companies, as long as they comply with applicable industry regulation. And it states that a disparate impact is not by itself sufficient to demonstrate an *intent* to discriminate.

Transparency

Like other recent AI laws, TRAIGA requires government agencies or healthcare providers to clearly disclose to individuals when they are interacting with a bot, “regardless of whether it would be obvious to a reasonable consumer that the consumer is interacting with an artificial intelligence system.” Disclosures must be clear and conspicuous, must be written in plain language, and may not use a dark pattern.

Sexually Explicit Content

TRAIGA also imposes restrictions on AI use in pornographic or sexual contexts, such as unlawful deepfakes, revenge porn, and exploitation of children.

Biometrics

Sometimes overlooked, TRAIGA also introduces amendments to Texas’ longstanding biometrics privacy law, the Capture or Use of Biometric Identifier Act (CUBI). Dormant for many years, CUBI has recently been enforced with vigor by Texas attorney general Ken Paxton, who secured fines north of a billion dollars under the law against Meta and Google. In the AI context, the amendments relax CUBI’s requirements, providing that individuals’ consent will not be required for “processing, or storage of biometric identifiers involved in artificial intelligence systems... unless performed for the purpose of uniquely identifying a specific individual.” This exemption would allow businesses to use biometrics for AI training without consent, including, for example, by web scraping, unless the data is used for identification purposes. At the same time, TRAIGA amends CUBI to state that “An individual is not considered to... have provided consent” to the processing of their biometrics “based solely upon the existence on the internet, or other publicly available source, of an image or other media containing one or more biometric identifiers.” In other words, only the voluntary public sharing of one’s image online would constitute CUBI consent.

Enforcement Mechanisms

TRAIGA comes with stiff penalties, and it’s enforced exclusively by the Texas attorney general. There’s no private right of action — although as previewed above, the law enables individuals to trigger the enforcement mechanism by filing complaints, and it requires the attorney general’s office to develop a reporting mechanism on its website to facilitate complaints of potential violations, similar to the one that exists under the TDPSA. The law provides businesses with a 60-day cure period. But after that, expect the attorney general to ask a court to impose civil penalties ranging from \$10,000 to \$12,000 per curable violation that wasn’t cured in time, or \$80,000 to \$200,000 per incurable violation — as well as an additional \$2,000 to \$40,000 per day for continuous infringements. Depending on the interpretation of what constitutes a single violation, these amounts can add up quickly, with fines for a *single violation* potentially soaring to millions of dollars per year.

TRAIGA isn’t all sticks; it also has carrots. These include an AI sandbox, which will enable participants “to obtain legal protection and limited access to the market in this state to test innovative artificial intelligence systems without obtaining a license, registration, or other regulatory authorization.” Once enrolled in the sandbox, a business is shielded from enforcement actions based on its activities as part of the program. The law states, “The attorney general 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.” Additional protections apply to businesses that demonstrate adherence to the NIST Artificial Intelligence Risk Management Framework.

TRAIGA also sets up a state Artificial Intelligence Council, comprising experts in law, ethics, and technology, who will be appointed by the governor, lieutenant governor, and speaker of the Texas House. The AI Council’s mandate includes ensuring

AI systems are ethical, are developed in the public's best interest, and do not harm public safety or undermine individual freedoms. Interestingly, TRAIGA also charges the council with "evaluat[ing] potential instances of regulatory capture, including undue influence by technology companies or disproportionate burdens on smaller innovators caused by the use of AI systems."

Summary

With its effective date looming, TRAIGA, the latest state AI law, which comes out of the red state of Texas, will occupy the minds of AI businesses across the country. In addition to strong civil rights protections and transparency obligations, TRAIGA embodies AI governance requirements that are likely to contribute to the development of a de facto national standard. Under it, the Texas attorney general is likely to become one of the most prominent AI

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

Omer Tene

Partner

otene@goodwinlaw.com

Boston | +1 617 570 1094