

At the Foot of the Rockies A New Approach to AI Policy Takes Shape



April 2026

WASHINGTON | CORE

Utah, a landlocked state in the American Mountain West, takes a light-touch approach to managing AI. With a policy of avoiding strict controls, making full use of AI for innovation, and keeping regulations to a minimum, the state aims to become a leader in AI implementation. The emergence of “Silicon Slopes” is taking place about an hour south of the state capital, Salt Lake City. The name of this AI corridor is derived from the slopes of the Rocky Mountains. Utah is navigating rapid population growth and rising demand for public services by using AI policy as a tool to address real-world challenges, from healthcare access to workforce shortages, positioning regulation not as a constraint, but as a driver of practical solutions. Washington CORE interviewed Jefferson Moss, Commissioner of the Utah Governor’s Office of Economic Development, to learn about the state’s AI strategy.

Building AI Policy Through Early Collaboration

Rather than waiting for AI to mature, Utah moved early to establish regulations that would support the industry’s growth. As generative AI began gaining global attention, the state engaged in extensive discussions to ensure that measures addressing risks such as fraud and discrimination would not come at the expense of innovation. The result was the enactment of the Artificial Intelligence Policy Act (AIPA), in March 2024, the first in the nation.

What sets Utah apart is that this policy did not emerge as a reactive response to AI’s rapid rise, but as a deliberate effort to get ahead of it. As tools like ChatGPT began to reshape public awareness, state leaders, led in part by then-legislator Jefferson Moss, convened a working group of industry experts, academics, and policymakers for months of discussion to map both the opportunities and risks of the technology before formalizing any regulatory approach. This early coordination allowed Utah to design a framework grounded not only in caution, but based on a clear understanding of how AI could be deployed in practice. Rather than asking how to restrict AI, the state focused on how to enable it responsibly from the outset.

A central objective of the policy was to provide regulatory certainty. As Moss noted, one of the greatest concerns for technology companies and investors is the risk of sudden regulatory shifts after significant capital has been committed. By clearly defining what is permitted and what is prohibited, Utah aims to create a stable environment in which companies can invest, build, and scale AI-driven solutions with confidence.

“Our top priority was to create an environment that maximizes the potential of AI while establishing clear guardrails for what is prohibited.” – Commissioner Jefferson Moss

Light-Touch Regulation Based on “Transparency,” “Accountability,” and “Balance”

The fundamental principle of the AIPA is encapsulated in the term “light-touch.”¹ Rather than stifling the industry with strict prohibitions, the philosophy is to ensure a minimum level of order while leaving as much room for innovation as possible. It is intentionally designed to prevent the government from overreacting to rapidly evolving AI technology.


The framework is built on three principles: “transparency,” “accountability,” and “balance.”

1. Ensuring “transparency” requires services using AI to explicitly disclose this fact to users. In regulated professions such as medicine, law, and accounting, the law prohibits AI from behaving like a human while concealing its true nature, thereby institutionally guaranteeing the foundation of trust with consumers.
2. Requiring “accountability” means even if AI provides misinformation and causes harm to consumers, the excuse that “it was the result of the AI’s independent judgment” will not be accepted. Businesses must treat AI as an integral part of their organization and bear full legal responsibility for its statements and actions.
3. By granting temporary exemptions in regulatory sandboxes for experimental development while concentrating regulations on high-risk applications such as hiring screening, healthcare, and credit assessments, the system achieves a “balance” that ensures safety and security without stifling technological progress.

A key feature of Utah’s AI regulations is that they are based on “actor regulation,” which focuses on the users who misuse AI rather than the tools themselves. The AI law first imposes minimum disclosure requirements. When individuals in professions requiring state licensing or certification interact with generative AI, they must explicitly state at the beginning of the interaction that the other party is an AI. Even in interactions with general consumers, if asked by the other party, they must disclose that the other party is a generative AI and not a human.²

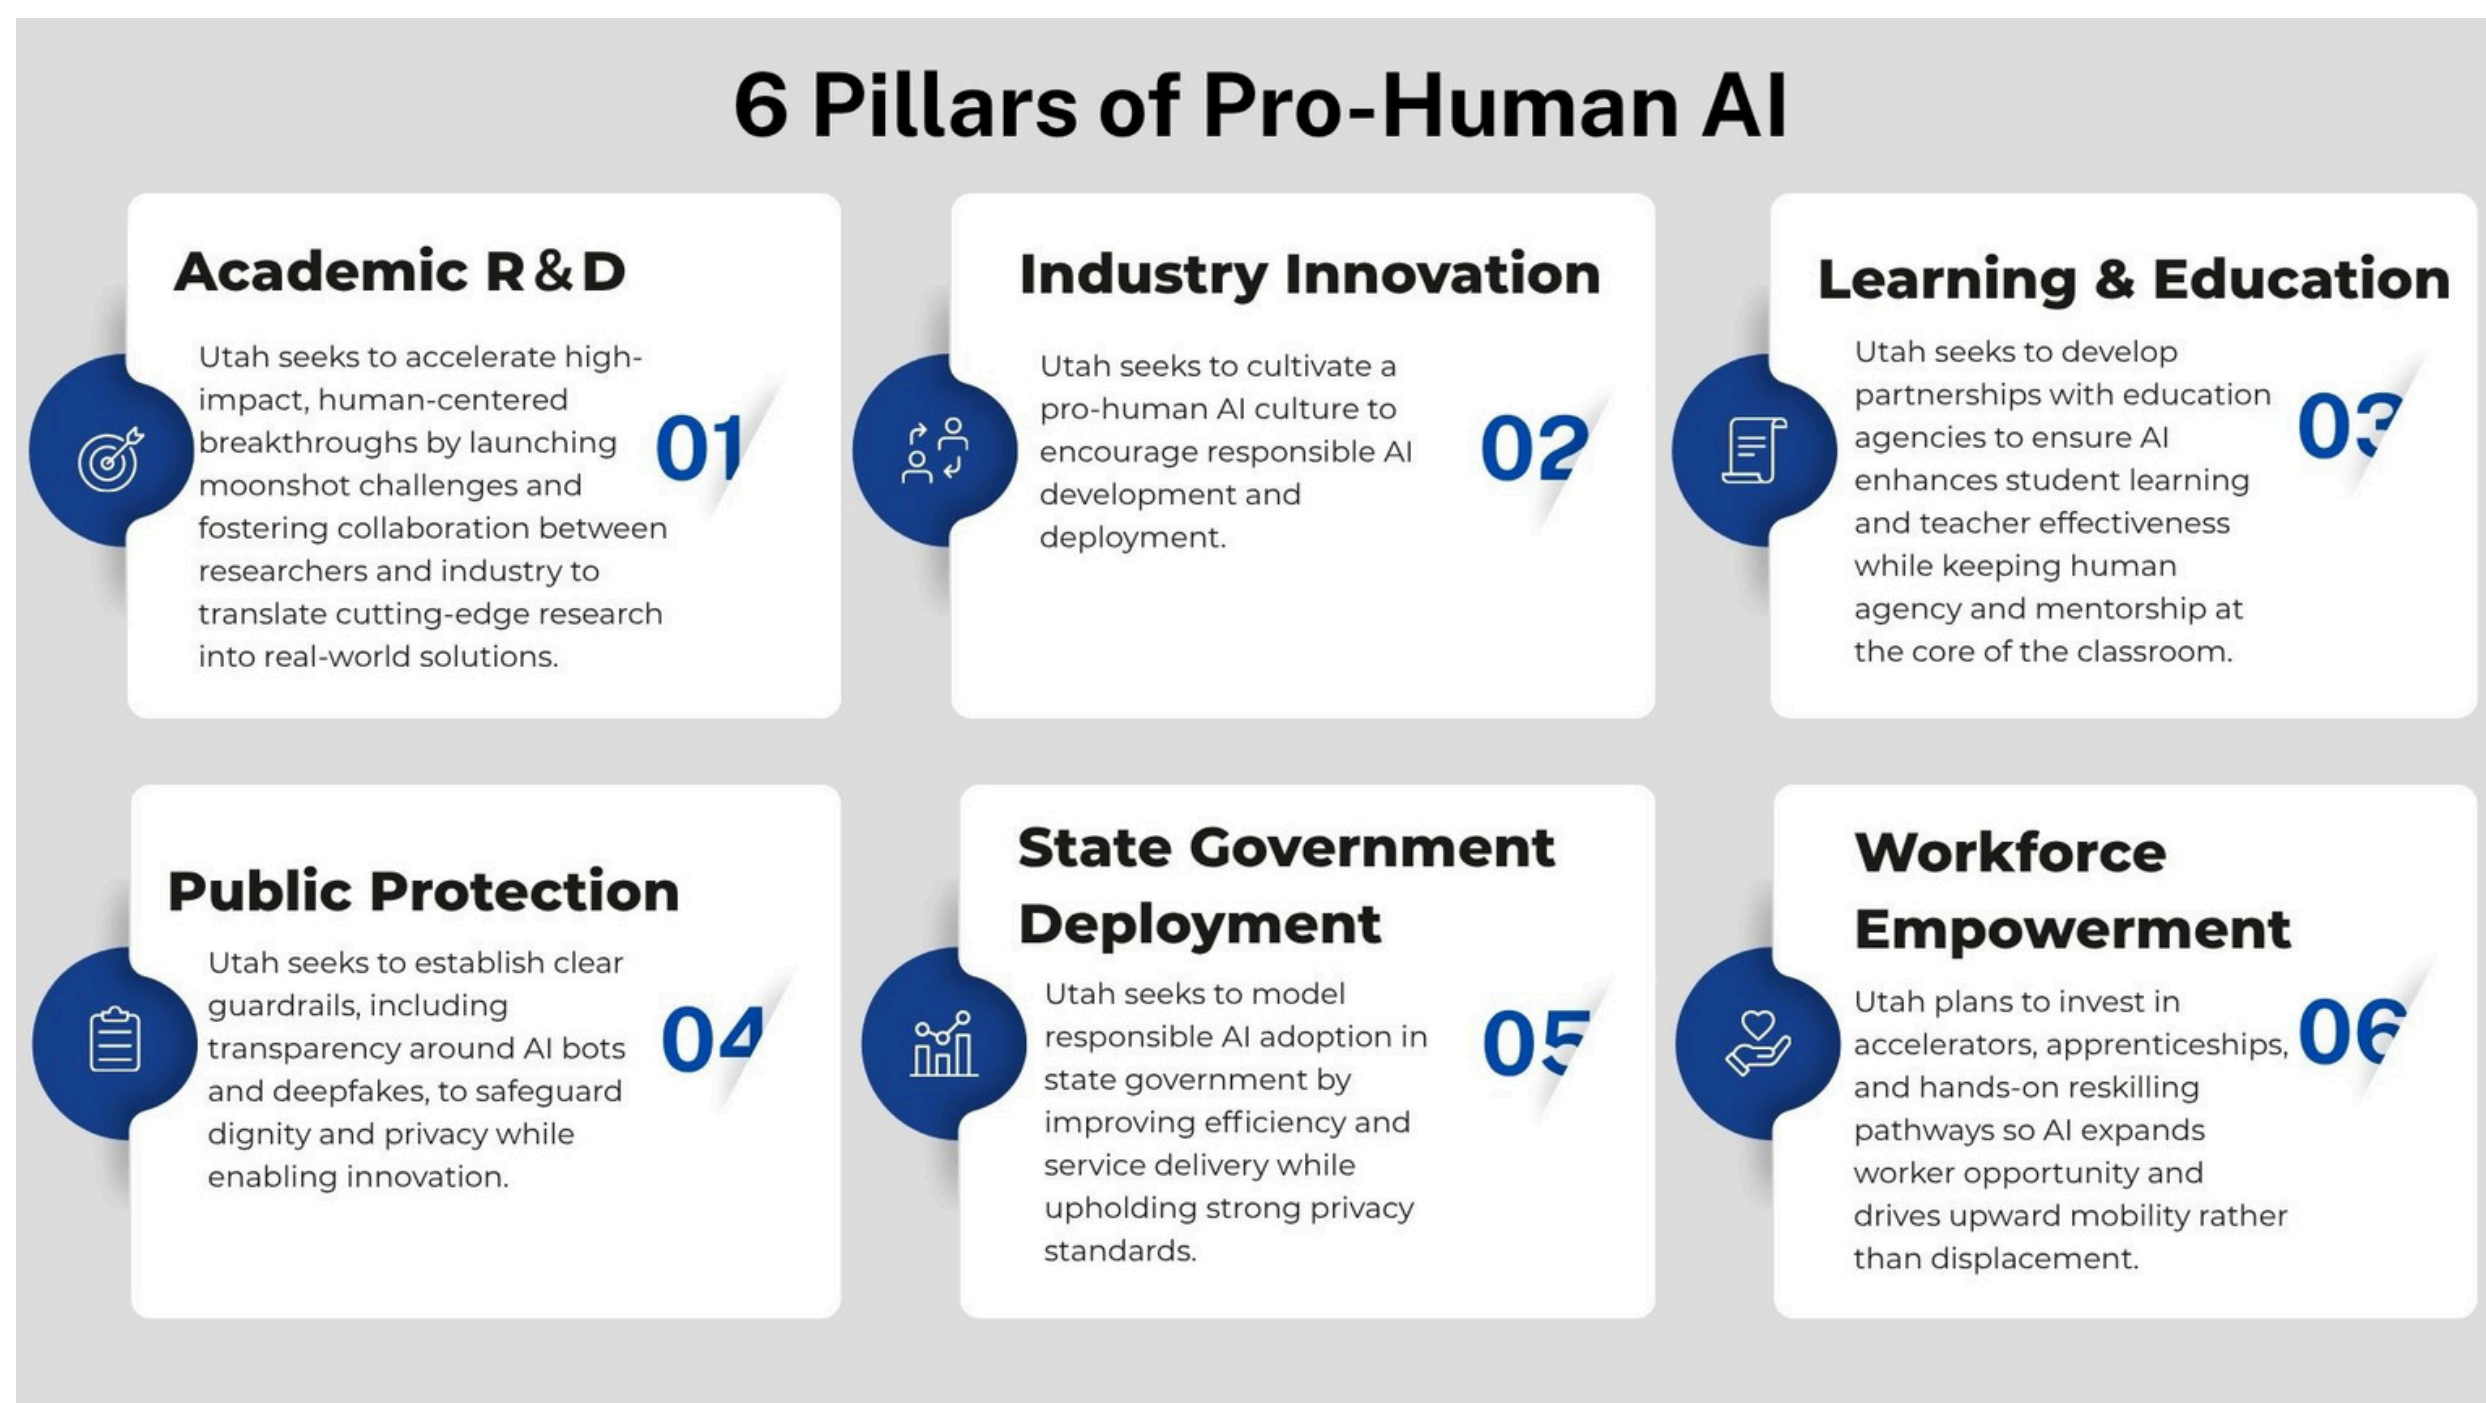
Utah is the first state in the U.S. to enact a consumer protection law that specifically regulates AI in this manner.³ Since then, the state has been progressively establishing a regulatory framework, including limiting AI disclosure obligations to “high-risk interactions” such as those involving health, finance, and biometric information; mandating disclosure that mental health chatbots are powered by AI and imposing advertising restrictions; strengthening regulations against identity impersonation, including deepfakes; and requiring transparency in police use of AI.

Regarding penalties for violations of these obligations, Utah sets a maximum fine of \$2,500 per violation, which is lower than in California (\$5,000) and Texas (up to \$200,000).⁴ Utah’s AI legislation is viewed as a simple, “light-touch” regulatory framework that offers predictability and a business friendly environment for companies and investors.



People First: "Pro-Human AI"

Following the enactment of the state law, Utah adopted the "Pro-Human Doctrine" and established a policy to maximize the use of AI for innovation. In December 2025, Governor Cox announced a \$10 million "Pro-Human AI" initiative at Utah's first AI summit.⁵ Guided by two core values, "AI must be human enabled" and "AI must be human enhancing," the initiative launched across six pillars: academic research, industrial innovation, education, public safety, state government AI deployment, and workforce development.



Source: Utah's Pro-Human AI Initiative⁶

Utah's "Pro-Human AI" strategy is supported by a set of coordinated initiatives designed to translate policy into practice. Among them is the state's "Moonshot" program, which promotes cross-sector collaboration between government, academia, and industry to identify priority areas for AI deployment and accelerate real-world solutions. Complementing this is the development of an "AI Learning Lab," which functions as a controlled environment for testing and evaluating AI applications before broader deployment.

The lab also serves as a real-time feedback mechanism, allowing technical experts to continuously assess emerging risks and opportunities, enabling policymakers to adapt alongside rapidly evolving technologies. Together, these efforts reinforce Utah's approach of treating AI governance as an iterative process, one that emphasizes experimentation, evidence gathering, and continuous refinement, rather than static rulemaking.

Moss describes the backdrop to the initiative as a response to an increasingly polarized debate around AI, often framed between extreme optimism (Boomers) and deep skepticism (Doomers). Utah, he explains, deliberately positioned itself between these two views, seeking to confront both the opportunities and risks of the technology head-on. Notably, this direction was not imposed from the top down.

Even before the governor's formal announcement, the state's Board of Higher Education had independently expressed support for a "Pro-Human" approach, while companies across the state were already moving in a similar direction. For Moss, this alignment across the public, private, and academic sectors reinforced a key conviction: Utah's AI strategy was emerging organically, rather than being driven solely by government mandate.

A critical pillar of this strategy is also workforce transformation. Recognizing that AI adoption will reshape labor markets, Utah has begun realigning its education and training systems to meet future demand. This includes efforts to expand AI-focused curricula, retrain existing workers, and better connect education pathways with industry needs. The emphasis is not only on preparing new talent, but also on ensuring that current workers can adapt, reflecting the state's broader "Pro-Human" philosophy that AI should augment, rather than displace, human capability.

A Regulatory Sandbox in Action

Utah has chosen the healthcare sector to become the primary testing ground for its "Pro-Human AI" philosophy. At its core is a simple question: can AI enhance human capabilities while leaving final decisions in human hands and can this balance work in real-world settings?

To answer this, Utah has approved four pilot projects that allow AI to be directly integrated into clinical and care settings under regulatory oversight. By allowing AI to take on specific, well-defined tasks, while maintaining human supervision, Utah is exploring whether AI can help extend the reach of limited medical resources and improve access to care without compromising safety.

This focus is driven by structural challenges within the state's healthcare system. Several regions have been designated as Health Professional Shortage Areas (HPSAs), and access to both primary and behavioral healthcare remains uneven.

An estimated 500,000 residents lack adequate access to mental health services, with some rural areas facing severe provider shortages. As Utah's population continues to grow, these pressures are expected to intensify.

Against this backdrop, AI is being positioned not as a replacement for healthcare professionals, but as a tool to expand capacity and improve access. The state's "Regulatory Relief" framework enables this approach by temporarily relaxing certain licensing constraints under controlled conditions, allowing new models of care to be tested while maintaining strict requirements for safety, privacy, and oversight.

The following pilot cases illustrate how regulatory flexibility is being applied in practice:

01 elizachat

Founded in Lehi, Utah, in 2023, ElizaChat develops AI powered mental health support apps for schools and is the first company with which the Office of Artificial Intelligence Policy (OAIP) has entered into a regulatory mitigation agreement (December 2024).⁷ Under Utah state law, a license as a mental health therapist was previously required if a conversation, ventured into therapeutic territory, even if it involved emotional counseling. Under this agreement, ElizaChat's AI chatbot is authorized to engage in conversations with teenage students in participating school districts regarding emotional processing, stress management, and goal setting.

However, if interactions cross into therapeutic practice, the company is required to address and correct the issue within 30 days. A safety management system is also in place requiring immediate notification to the school's emergency contacts and connection to a licensed therapist if serious cases, such as signs of self-harm or abuse.

The system is designed to function as a "gateway" to care rather than a substitute for therapists, serving to address the state's shortage of school counselors. ElizaChat is currently expanding into school districts in five states, including Illinois and Oklahoma.

02 doctronic

Headquartered in New York, Doctronic operates a 24/7 AI-native health platform across all 50 U.S. states. The company signed a regulatory agreement with Utah in December 2025, becoming the first in the U.S. to be legally authorized to use AI to assist with prescription renewals.

Traditionally, only licensed physicians could renew prescriptions, but under this agreement, Doctronic's AI is authorized to handle the prescription renewal process for 192 types of low-risk medications, including treatments for chronic conditions such as high blood pressure and cholesterol, without requiring direct physician oversight. Complex cases are escalated to physicians on a 24/7 basis, and health data privacy laws and cybersecurity standards continue to apply.

As a result of advancing pilot programs and service deployment under this framework, the company is rapidly expanding its user base and business scale. The company has achieved significant success in fundraising, growing 15-fold in less than six months since completing its Series A round and reaching eight-figure annual revenue. It is reported that weekly users exceed 300,000, and the repeat patient rate has tripled.⁸

03 dentacor

Founded in Utah in 2021, Dentacor provides mobile dental hygiene services to shelters, recovery facilities, and transitional housing. The company has secured a 12-month regulatory relief agreement from the OAIP.⁹ Under Utah's dental laws, dental diagnoses based on X-ray images could previously only be made under the direct supervision of a dentist.

However, under this agreement, dental hygienists can now diagnose specific dental conditions using AI-assisted radiographic diagnostic tools.

However, the diagnosis is only confirmed if the results from both the hygienist and the AI match. If there is a discrepancy, the case is escalated to a dentist. The scope of eligible procedures is limited, and there is an obligation to obtain informed consent from all patients and report to the OAIP monthly. This initiative is meant to simultaneously advance the social mission of bringing access to dental care to those furthest from it, such as the homeless and low-income populations, and to verify the safety of AI-based diagnosis.

04 legion HEALTH

Founded in 2021, Legion Health, a provider of online psychiatric services, was authorized in April 2026 under Utah's regulatory sandbox to allow its AI chatbot to autonomously renew prescriptions for psychiatric maintenance medications, typically low-dose antipsychotics or antidepressants used to prevent relapse and maintain patient stability, without requiring prior physician approval.¹⁰

This is the first instance in the world where a government has granted an AI system the authority to autonomously prescribe psychiatric medications.

The scope is limited to 15 types of low-risk maintenance medications, and the system is not permitted to initiate new prescriptions, adjust dosages, or prescribe controlled substances. As a safety measure, a three-phase rollout has been designed: the first 250 cases will be directly supervised by a physician, the next 1,000 will undergo post-hoc review, and only after achieving an approval rate of 98% or higher will the system transition to autonomous AI operation. Offered as a \$19 monthly subscription service, the program aims to provide stable patients living in rural areas or counties without psychiatrists with a means for rapid prescription renewals.

The “Utah Model”: Global Attention and Future Challenges

Utah’s “Human-Centered AI Doctrine,” which seeks to minimize regulations and harness AI for innovation, has attracted attention both domestically and internationally. As the U.S. Congress intensifies its efforts to nullify state AI regulations, Utah is positioned as one of the leading states, alongside New Hampshire, New York, and North Carolina, in spearheading an open letter to Congress.¹¹ According to Commissioner Moss, a UK government official responsible for AI policy dubbed this approach the “Utah Model.” Upon hearing this, Moss reportedly smiled to himself, thinking, “That’s exactly what we’re aiming for.”

Utah isn’t solving everything. Moss identified three challenges. First of which is the anxiety and resistance toward AI among teachers and the public. Careful communication is essential to address concerns that jobs will be replaced. Second is the lack of awareness regarding support measures, specifically that robust reskilling programs are not reaching those who need them most. Third is the issue of the digital divide, where the gap between those with and without AI skills is rapidly widening in rural areas.

Under the banner of ‘Pro-Human AI,’ Utah is rolling out evidence-based initiatives, including the aforementioned regulatory relief, at a rapid pace. At the same time, the state is intentionally avoiding a one-time policy approach, instead creating a system that is continuously reassessed and refined based on emerging needs, risks, and opportunities.

Looking ahead, Utah is working to build an environment in which AI governance evolves alongside the technology, with regular evaluation of what challenges matter most and where intervention is most effective.

“Long-term, we want to have the most educated workforce and the best economy. We want to incentivize all those factors. And if, five years from now, some of these initiatives like the moonshots create something that has a massive impact on the world—and we can say Utah helped lead that—those are the things we really want to lean into.” - Commissioner Jefferson Moss

As population growth continues to place pressure on public systems, Utah is positioning AI policy as a tool for managing real-world demands, rather than simply regulating emerging technology. While its approach is still evolving, early efforts suggest a model of governance built on experimentation, trust, and adaptability—one that other regions may increasingly look to as they navigate the same challenges.

Source: Visit Utah¹²



Endnotes

- [1] <https://www.techbuzznews.com/utah-leads-on-ai-legislation/>
- [2] <https://le.utah.gov/~2024/bills/sbillenr/SB0149.pdf?utm>
- [3] <https://cdp.cooley.com/utah-colorado-pave-way-for-ai-specific-state-laws-is-your-company-ready-for-the-impending-regulation-wave/>
- [4] <https://www.skadden.com/insights/publications/2024/04/utah-becomes-first-state>
- [5] <https://www.cdomagazine.tech/aiml/utah-launches-10-million-pro-human-ai-initiative-to-prep-workforce-and-expand-adoption>
- [6] <https://prohumanai.utah.gov/six-pillars/>
- [7] <https://commerce.utah.gov/2024/12/02/news-release-utah-department-of-commerces-office-of-artificial-intelligence-announces-first-regulatory-mitigation-agreement/?utm>
- [8] <https://www.businesswire.com/news/home/20260324814372/en/Doctronic-Raises-%2440M-Series-B-Following-Breakthrough-as-First-AI-to-Legally-Renew-Prescriptions-in-the-U.S.?utm>
- [9] <https://commerce.utah.gov/ai/agreements/dentacor/>
- [10] <https://www.pcmag.com/news/ai-can-now-prescribe-you-psychiatric-medication-in-utah>
- [11] <https://www.naag.org/policy-letter/state-attorneys-general-urge-congress-to-preserve-local-authority-on-ai-regulation/>
- [12] <https://www.visitutah.com/places-to-go/cities-and-towns>

Interviewee

Jefferson Moss, Commissioner, Governor's Office of Economic Development (GOED)

He has a diverse background as an entrepreneur, venture capitalist, and legislator. After working in private banking at Credit Suisse and KeyBank, he founded several startups. He has served as Chair of the Endowment Investment Committee at Utah Valley University (UVU), founder of the UVU Wolverine Fund and the Utah Innovation Fund, Vice Chair for Innovation on the Utah State Board of Higher Education, member of the Utah State Board of Education, Saratoga Springs City Councilmember and member of the City Planning Commission, Executive Director of The Point Innovation District, and Majority Leader of the Utah State House of Representatives.



WASHINGTON | CORE

Washington CORE, L.L.C. is an independent consulting & research firm providing strategic research, analysis, and advisory services. Founded in 1995, Washington CORE leverages in-depth research capabilities coupled with extensive global networks in both the public and private sectors, to deliver clarity and insight to prepare our clients for success in an ever-changing global landscape.

Please visit <https://www.wcore.com> for more information.

<https://www.wcore.com>