

Ai4 Conference: AI Politics and Policy

August 4-6, 2026

Venetian Expo Center, Las Vegas, Nevada

The current politics and future policy surrounding AI was a major topic of discussion at the Ai4 conference. To start off the discussion, Andrew Yang, former presidential candidate in 2020 and CEO of Nobel Mobile, framed the politics of AI as a widening divide between technological optimism and public distrust, driven largely by economic insecurity and job displacement. He argued that many Americans, especially recent college graduates who are avid AI users, see AI as a threat because they already struggle to find stable work, and large companies are using automation to justify layoffs. Communities resist data-center construction because they perceive the benefits flowing to corporations while the costs of these centers, such as water use and noise, fall on local residents.

Yang warned that AI value is increasingly “sucked into the cloud,” contributing to job polarization in which high- and low-income roles grow while the middle disappears. He called for a human-centered approach to capitalism that protects workers, limits digital harms, and addresses the anxiety caused by social media and technological change. Politically, Yang argued that the United States needs a new movement beyond traditional parties, which is one capable of forging a “grand bargain” around AI, economic security, and community well-being. He hopes to advance these principles through the Forward Party.

The Historical Context and Future of AI

Geoffrey Hinton, considered the godfather of AI; Fei Fei Li, considered the godmother of AI and the founder of World Labs; and Andrew Ng, founder of Deep Learning AI, discussed the historical context and the future of AI.

Geoffrey Hinton warned that AI systems are beginning to behave in ways programmers did not intend, noting that AI is doing things that the programmers did not intend it to do, including both routine and egregious actions. He argued that large models increasingly look around for security flaws and sometimes refuse instructions or generate original proofs. Hinton believes many routine intellectual jobs will disappear. He asked: “Will people whose jobs are destroyed be retrained for new jobs?” He sees promise in AI tutors that could eventually match human teaching, but he insists regulation is essential. Regulation should be the steering wheel, not the brakes. He also cautioned that open-weight models make misuse easier because the genie is out of the bottle. He predicted that within five years either AI tremendously improves everyone’s life or a rogue AI has done a cyber-attack.

Fei-Fei Li urged a human-centered approach, arguing that public discourse is distorted by extremes. Instead, the public needs education to learn how to use AI as a tool. She emphasized that jobs consist of multiple tasks, and AI will reshape some but not all of them. Fear-based

rhetoric does not help people to learn. Instead, society must invest in reskilling, community support, and motivation: “Education is self-agency.” She warned that productivity gains do not automatically create shared prosperity and called for updated regulatory frameworks that protect people while enabling innovation. Li rejected the false dichotomy between open and closed systems: “We can have both open source and closed models.” The use of open and closed systems should depend on context and used responsibly.

Andrew Ng, founder of Deep Learning AI, pushed back against claims of imminent mass unemployment, noting that only 1.4% of layoffs are because of AI. Instead, company layoffs are because of over-employment. Companies are letting go of employees who are not needed anymore using AI as their excuse. He argued that workers should gain new skills and broaden their scope because companies will need employees who are more generalists, who can look at the whole picture and problem solve creatively. Ng reassured students that humans retain a fundamental content advantage over AI because we are better at understanding the world. He emphasized that AI should empower people rather than replace them. Ng envisions a future where AI helps eradicate illiteracy, cure disease, and expand global access to intelligence.

Developing Trust in AI

Companies are struggling to build trust in AI because developers, businesses, and customers all worry about reliability, security, and unintended behavior. As AI systems become more capable and self-updating, people fear losing control over the software that runs their products. Trust grows only when AI reduces friction, improves affordability, and works consistently across devices. Businesses must show clear value—faster processes, easier authentication, and simpler user experiences—while maintaining strong governance and security. Developers need transparent systems they can evaluate, and customers need confidence that AI will behave safely. Ultimately, trust in AI emerges when companies, engineers, and users all participate in shaping systems that are reliable, accountable, and aligned with human oversight.

Government Uses of AI

Government officials described how AI is becoming essential for managing rising workloads, modernizing outdated systems, and improving public service delivery. Roman Jankowski, Federal Dept. of Security, explained that agencies must respond to transparency requests within 20 business days, and AI now helps process massive volumes of documents, automate routine interactions, and re-engineer legacy systems that previously relied on paper. This frees up staff to focus on higher-value tasks while keeping humans in the loop for final decisions.

Robert Fulk, Indiana state government, emphasized that AI reduces costs and accelerates services in areas such as business regulation, immigration processing, and privacy requests. Indiana uses a mix of hyperscaler tools and custom in-house systems to meet strict privacy requirements, ensuring that humans—not AI—make determinations in high-stakes cases. He stressed that vendors must understand government problems and deliver solutions that are easy to implement.

Christian Napier, Utah state government, highlighted Utah's "pro-human AI" initiative, which uses AI to enhance—not replace—state employees. After early failures with generic chatbots, Utah built specialized agents that analysts trust for contract review and compliance checks. AI increases capacity for overburdened agencies, allowing employees to focus on complex work. Napier noted that states rely on a broad ecosystem of partners and must set guardrails before deploying synthetic hosts or chatbots to the public. Across all speakers, the theme was clear: AI is becoming a critical tool for government, but human oversight, privacy protection, and trustworthy partnerships remain central.

AI and Midterm Elections

The political landscape heading into the midterm elections is being reshaped by public skepticism toward AI, intense industry lobbying, and growing conflict over data-center expansion. Polling shows that concern about AI now outweighs excitement by roughly three to one, and only the wealthiest Americans expect AI to improve their lives. Job-loss anxiety crosses party lines, with more than half of voters worried about economic disruption. This skepticism is becoming a central campaign theme, especially in communities where large data-center projects provoke fears about water use, noise, land impact, and the sense that "someone else is making money but not you." Adoption of AI tools does not equal consent, and many voters feel the technology is being pushed on them without a clear national strategy.

Washington's response has been fragmented. Congress has introduced many AI-related bills, but almost none have passed, leaving what analysts describe as "a security posture, not a strategy." The executive branch has issued multiple orders, but agencies are still reacting to incidents such as recent agent-hacking events involving major AI companies. In this vacuum, alliances, advocacy groups, and state governments are beginning to shape their own rules, creating a patchwork of oversight.

The midterms are likely to produce divided government, with analysts expecting one chamber to flip while the other remains under the opposite party's control. Comprehensive federal AI legislation is unlikely, but more oversight, subpoenas, and investigations into AI companies are expected. Campaign ads increasingly feature AI as a wedge issue—especially around data centers, corporate power, and community impact. After the election, U.S.–China AI policy will become a major focal point, and observers anticipate a "catalyzing shock"—either a breakthrough or a crisis—that could reset the national debate. The core question remains: what should government do when a handful of firms hold such concentrated technological and economic power?

The Governance Imperative for AI

AI governance has become both urgent and difficult because technology evolves faster than legal and regulatory systems can respond. Policy deadlines pass, executive orders shift, and new incidents, such as models recently escaping sandboxes or unexpected agentic behavior, highlight how quickly risks emerge. The core challenge is deciding *what* to govern because

different AI tasks create different risks, and each requires a distinct regulatory approach. Malfunctions, liability, data-center impacts, geopolitical tensions, and the rise of agentic AI all demand separate solutions. AI governance is hard because AI is constantly changing, and regulators must constantly update their frameworks.

Four interlocking tensions shape the governance landscape: (1) U.S.–China competition, (2) open versus closed models, (3) freedom versus authoritarian control, and (4) public versus private power. Chinese models are paradoxically among the most open, while U.S. federal policy avoids banning open models but regulates closed ones more heavily. In practice, governance is emerging through a patchwork of actors, such as investors inserting governance terms into deals, insurers requiring standards, auditors developing methods to evaluate black-box systems, and states writing their own AI laws. Federal executive orders remain voluntary, not command-and-control, leaving gaps that states are attempting to fill.

This fragmentation raises questions about preemption, innovation, and cost. National companies struggle with inconsistent state rules, and small entrepreneurs—especially in Europe—face heavy compliance burdens. Experts argue that the federal government must set baseline regulations while allowing states flexibility for local governance. Without federal clarity, standards may become the default tool for shaping AI behavior.

The political reality is that no major AI legislation is likely to pass before the mid-term election, leaving a vacuum in which oversight, subpoenas, and investigations will grow. The imperative is clear that governance must balance innovation with safety, define liability, protect the public, and ensure that AI develops within a coherent national framework rather than through reactive, fragmented rules.

AI on the ballot

Community reactions to AI are increasingly driven by the visible expansion of data centers, which have become the most tangible symbol of AI technology. Kevin Frazier, a law professor at UT Austin, explained that voters often cannot connect abstract AI benefits to their daily lives, but they clearly understand the physical impacts of massive data-center projects and their impact on noise, water use, energy demand, and land disruption in their communities. As a result, data centers have become a hot ballot issue, even though many residents are not opposed to AI technology itself. Instead, they are uncertain about who pays for the electricity, who receives the tax breaks, and whether local communities gain anything from these facilities.

Counties face antiquated laws and uneven negotiation power with hyperscalers, giving hyperscalers the upper hand in building these large centers. On the other hand, voters worry about AI affecting job displacement, especially for young people entering the workforce. Construction jobs exist with data centers, but long-term employment is limited, fueling concerns that the jobs are gone once the building is done. Wealth concentration also shapes public sentiment, with frustration that four major companies are making all the decisions while ordinary residents see few benefits.

Frazier noted that AI is becoming a political issue similar to COVID because it changes rapidly, is confusing, and lacks a clear playbook. Many voters do not know where candidates stand, and NIMBY (Not In My Back Yard) reactions intensify as communities struggle to discern what is good or harmful. AI's presence on the ballot reflects this uncertainty and the need for clearer public engagement.

The Need for Stronger Federal AI Governance

Federal efforts to govern AI are falling short, leaving businesses and regulators struggling with uncertainty. Experts argue that Washington cannot simply say it will do nothing for ten years, because AI is already more complex and potentially dangerous than past technologies, such as aviation. Congress often misunderstands what AI is and how existing laws already apply, while states create their own rules, producing a patchwork that burdens interstate commerce. Over-regulation at the state level could reduce adoption and harm small businesses, yet full federal preemption could leave industry unrestrained and uninsurable, especially for agentic AI systems that insurers cannot currently assess.

The core problem is institutional weakness because agencies lack the expertise and authority to enforce existing laws, and political backlash pressures lawmakers into rushed or symbolic actions. Experts emphasize that the United States needs stronger federal regulation, clearer standards, and better-resourced institutions to manage AI's rapid evolution while preserving innovation and economic competitiveness.

AI Geopolitical Competition

The United States is locked in an intense AI race with China, but policymakers increasingly recognize that competition alone cannot ensure safety or global leadership. Experts argue that the U.S. must help build international technical standards, which are distinct from regulation, because “everyone agrees we have an evaluation gap” and cannot yet reliably assess AI systems. Without trusted standards, the U.S. risks losing influence, much as it has in other technologies where China now shapes emerging-market adoption. Europe, meanwhile, is moving ahead with strict regulatory frameworks such as the EU AI Act, giving it outsized power to define compliance norms even as the U.S. favors lighter regulation. The ideal future resembles global aviation safety: shared international standards that ensure trust, reduce risk, and allow both proprietary and open-source models to coexist. Achieving this requires technical cooperation, political alignment, and renewed U.S. engagement on the world stage, especially as open-source models—many from China—accelerate worldwide adoption.

An AI Human-centered Focus from Policy to Practice

A human-centered AI strategy requires institutions to move from abstract policy debates to practical, transparent, and inclusive implementation. The University of Michigan's approach shows that responsible AI begins with a culture that listens to all voices, including to give a voice to the silence, and ensuring that curiosity guides experimentation. Human judgment must

remain central, supported by intentional use of AI tools and clear understanding of when *not* to use them. Responsible practice also demands a vendor-agnostic architecture, sovereign data strategies, and tools built around the needs of the people who do the work. AI must be shaped, rather than simply adopted, through ethics, core values, and thoughtful AI governance as a part of the process from the start. As agentic AI scales, organizations must integrate transparency, critical thinking, and environmental awareness into their AIQ. The path from policy to practice should include the building of systems for humanity, protecting human agency, and designing AI that is ready for tomorrow while grounded in the responsibilities of today.