HARVARD Kennedy School
MOSSAVAR-RAHMANI CENTER
for Business and Government

The People Have Spoken: The Tech Industry, Civil Society, and the U.S. Artificial Intelligence Action Plan

Paulo Carvão
Harvard Kennedy School

Mizuki Yashiro
Harvard University

Shaurya Jeloka
Harvard University

June 2025

M-RCBG Associate Working Paper Series | No. 261

Mossavar-Rahmani Center for Business & Government
Weil Hall | Harvard Kennedy School | www.hks.harvard.edu/mrcbg

The People Have Spoken

The Tech Industry, Civil Society, and the U.S. Artificial Intelligence Action Plan

Paulo Carvão, Mizuki Yashiro, Shaurya Jeloka

Abstract

In early 2025, the White House issued a Request for Information to guide the development of the national AI Action Plan. The responses submitted by entities ranging from Big Tech and venture capital (VC) firms to universities and civil society organizations provide an unprecedented snapshot of stakeholder expectations for U.S. AI policy. This paper examines the positions articulated in these submissions through a mixed-methods analysis combining qualitative review and natural language processing. We identified patterns of alignment and divergence across stakeholder groups and then mapped them onto a previously developed typology of AI worldviews. Our findings reveal that while the private sector broadly supports American leadership and light-touch regulation, it is far from monolithic. Civil society emphasizes the importance of risk mitigation, trust, and accountability. Understanding these competing perspectives is critical as the federal government finalizes its AI Action Plan and determines how to address the tension between innovation and public interest concerns.

Keywords

Artificial Intelligence; Policy; AI action plan; AI governance; U.S. Government; Tech Industry

Biographical Notes

Paulo Carvão is a Senior Fellow at Harvard Kennedy School's Mossavar-Rahmani Center for Business and Government, specializing in Tech Policy and AI. He is a former IBM executive with three decades of experience in digital transformation, cloud, and AI integration. Currently an advisor to tech startups and a venture capital investor, Paulo has held fellowships at Harvard's Advanced Leadership Initiative and Safra Center for Ethics. He serves as the inaugural Entrepreneur-in-Residence at the Harvard Kennedy School's GovLab.

Mizuki Yashiro is an A.B. candidate in Economics and Government on a Data Science Track at Harvard. She has experience working within the fields of venture capital and private equity, with Paradosi Partners, Xfund, and the Phoenix Fund, and is the Co-Founder of Admitvantage, a college consulting educational startup. At Harvard, she serves as Deputy Chair of the Institute of Politics GovLab and as the Director of Strategy at the Harvard Venture Capital Group, where she leads diligence teams and investment reporting. Mizuki's experience spans private equity, consulting,

and applied AI governance research. She is passionate about the intersection of emerging technology, venture capital, and civic infrastructure.

Shaurya Jeloka is an A.B. candidate in Computer Science and Economics at Harvard. His work spans intelligent algorithms for robotic systems, predictive modeling for M&A risk mitigation, and performance analysis tools used across academic and industry contexts. His industry experience includes work at Amazon Robotics, where he developed tooling to improve the performance and reliability of robotic systems. At Harvard, he works as a research assistant on foundational AI research and as a project lead at the Undergraduate Machine Intelligence Club, where he contributes to AI tools that support legislative research and policy analysis.

Table of Contents

Table of Figures

Introduction

In January 2025, the U.S. government issued an Executive Order calling for the removal of barriers to American leadership in artificial intelligence (AI). (The White House, 2025a) This directive included the development of an Action Plan within 180 days. In February, the White House Office of Science and Technology Policy (OSTP) issued a Request for Information (RFI) to support designing the plan. Academia, industry associations, the private sector, non-governmental organizations, and citizens answered the call. By late April, the OSTP published 10,068 comments.

This overwhelming response created a unique time capsule of public opinion. Our piece unpacks this trove of insights and highlights a gap between the perspectives of the tech industry and civil society. They offer fundamentally different governance visions: industry prioritizes innovation, global competitiveness, and regulatory preemption, while civil society emphasizes equity, accountability, and democratic safeguards. Bridging this divide is essential for designing a framework that stimulates innovation while earning public trust.

The industry does not speak in unison. Instead, their input includes multiple voices reflecting diverse interests and business strategies. One distinction emerged between Big and Small Tech, and within each group, further variation reflects the categories discussed in "Governance at a Crossroads: Artificial Intelligence and the Future of Innovation in America" (Carvao et al., 2025): accelerationists, responsible AI advocates, open AI innovators, safety advocates, promoters of public interest AI, and national security hawks. This framework is used throughout this work.

The industry is driven by American leadership, national security, and innovation. Leadership is often framed in terms of export controls, infrastructure investments, and a push for the diffusion of American technology. Within Big Tech, two camps emerge: traditional enterprise technology providers on one side, and, on the other, large frontier labs and internet platforms. Beyond the shared emphasis on leadership, national security, and innovation, traditional providers frequently position themselves as responsible AI advocates, whereas frontier labs and internet platforms tend to be accelerationists. VC firms also exhibit strong accelerationist behavior. Advocacy for safety is largely absent, with Anthropic representing a notable exception. Overall, the industry favors minimal regulatory intervention.

Civil society, by contrast, emphasizes trust in AI systems as a precondition for widespread adoption. Their concerns center on ensuring AI safety, mitigating harms, protecting civil liberties, and addressing copyright and intellectual property issues. These stakeholders remind us of the importance of investing in education and workforce development and are generally receptive to governance structures aimed at mitigating AI risk.

Research Questions and Approach

To support the development of an AI Action Plan, the White House OSTP, in coordination with the Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO) and the National Science Foundation (NSF), issued an RFI on

February 6, 2025, which received comments until March 15. (National Science Foundation, 2025) On April 24, the White House and NITRD published 10,068 responses. (Networking and Information Technology Research and Development, 2025; The White House, 2025b)

This unprecedented volume of inputs offers a representative snapshot of America's opinions on AI. Each respondent sought to shape federal policy in a direction aligned with their values or commercial interests. We posed the following questions:

- What can be learned from the responses, and what main themes emerge?
- What is the industry advocating for, and does it speak with one voice?
- Is Big Tech's position aligned with the VC community and Small Tech?
- Do the industry responses match the patterns previously identified (accelerationists, responsible AI advocates, open AI innovators, safety advocates, promoters of public interest AI, and national security hawks)?
- What positions have been expressed by civil society and universities?
- Does industry align with civil society?

To proceed with these questions, we employed two complementary methodologies: qualitative analysis of selected submissions and quantitative analysis of the complete dataset from NITRD. (Networking and Information Technology Research and Development, 2025)

An in-depth **qualitative review** was performed with forty-five (45) responses, in six groups:

1. Big Tech (16) – Adobe, Amazon, Anthropic, CISCO, Dell, Google, IBM, Intel, Lockheed Martin, Meta, Microsoft, OpenAI, Qualcomm, Salesforce, Uber, Verizon.
2. Small Tech (5) - Cloudflare, CrowdStrike, ServiceNow, Snowflake, Workday.
3. Venture Capital (4) – Andreessen Horowitz, General Catalyst Institute, National Venture Capital Association, True Ventures.
4. Industry Associations (6) – AI Alliance, Business Roundtable, Business Software Alliance, Chamber of Progress, Global Data Alliance, U.S. Chamber of Commerce.
5. Civil Society - other companies and corporations (3): Johnson & Johnson (J & J), JP Morgan Chase (JPMC), PepsiCo.
6. Civil Society- academia, advocacy associations, and think tanks (11) - Association of American Universities, Association for Computing Machinery, Association of Public and Land-grant Universities, Association of University Presses, Center for AI and Digital Policy, Center for AI Policy, Electronic Frontier Foundation, GETTing-Plurality network at Harvard Kennedy School of Government, International Center for Law & Economics, IEEE, The Leadership Conference on Civil and Human Rights.

The classification of organizations was based on public information and our informed judgment.

Big Tech was defined expansively to include not only the traditional "Magnificent Seven," Alphabet (Google), Amazon, Apple, Meta (formerly Facebook), Microsoft, NVIDIA, and Tesla

(Tesla has not submitted an RFI response), but also other large-cap technology firms and the two largest AI frontier laboratories (Anthropic and OpenAI).

All other technology companies not included in this list of sixteen were considered Small Tech. Moreover, we included the responses of venture capitalists in the Small Tech analysis.

The RFI responses included inputs from industry associations and trade groups. To better understand industry voices not captured in the previous three categories (Big Tech, Small Tech, VCs), we selected six associations, mainly representing companies instead of unions or trade groups.

Other companies or corporations not primarily in the technology sector were placed into the Civil Society category. In our final analysis, the Civil Society group included academic and research institutions, professional associations, advocacy organizations, think tanks, non-tech companies, industry associations, and trade groups.

In addition to the qualitative analysis, the 10,068 responses were reviewed using Machine Learning (ML) techniques. This approach allowed us to identify thematic clusters and frequently cited regulatory or ethical concerns across the diverse stakeholder groups.

Data and Methods

Classification Pipeline

We used large language models (LLMs) to classify each comment along two dimensions: institution type and AI worldview. We performed initial cross-model validation by comparing the classifications generated by OpenAI's GPT-4 Turbo and Google's Gemini 1.5 Flash. The models were accessed via their API interface and prompted to parse the responses to extract the following information:

- Submitter name and type
- Institution type (big tech, small tech, civil society)
- Specific organizational tag (startup, nonprofit, government, media, etc.)
- Two AI worldview labels (from a set of six possible options)
- Justifications for the institutional and AI worldview classifications
- A filter tag indicating whether the comment should be retained or discarded

Our full set of AI worldviews, as described in the introduction, is as follows: Accelerationists, Responsible AI Advocates, Open AI Innovators, Safety Advocates, Public Interest AI, and National Security Hawks. Since the textual analysis revealed that many comments held multifaceted views, we prompted the model to assign two worldview classifications (primary and secondary) to each comment.

Upon comparing the two models, our results showed nearly perfect agreement on institution type (≈99%) and majority agreement on AI worldview classification (≈85%). Given Gemini's significantly lower costs, we selected Gemini 1.5 Flash for full-scale analysis over GPT-4. To minimize classification errors and address category overlap (particularly between Responsible AI Advocates, Public interest AI, and National Security Hawks), we refined the prompt with a set of counterexamples and specific rules for ambiguous areas. Furthermore, to ensure that only substantive and policy-relevant comments were included in our analysis, we asked Gemini to flag and discard comments that lacked meaningful policy input. This process filtered out 485 responses, targeting comments such as *"Please bring some data centers to Utah!"* or *"AI has no place in the future of America. It is a huckster scam…"* We are not commenting on these views themselves—they were filtered out simply because they did not provide constructive or actionable input, rendering them ill-suited for our goals. Our final prompt is as follows:

```
You are an AI assistant tasked with analyzing public comments submitted to the U.S.
government's AI Action Plan Request for Information (RFI) in 2025. Your goal is to
extract structured data from these comments and categorize them based on their
submitter and AI-related views.

Important: Do not rely solely on key phrases. Always infer the submitter's underlying
priorities or concerns. Words like "governance", "risk", or "transparency" can appear
in different worldviews -- focus on why they are used, and what problem the comment is
trying to address.

You may use the filename (which may contain the submitter's name or organization) to
help determine the submitter identity.

Your task is to extract the following information from each comment:

1. submitter_name: The name of the person or organization who submitted the comment.
2. submitter_type: Is this an "individual" or "institution"?
3. main_group: One of: "Big Tech", "Small Tech", or "Civil Society". Use the following
rules:

  - "Big Tech": Only include the following companies and their direct representatives:
OpenAI, Microsoft, Meta, Intel, IBM, Google, Anthropic, Amazon, Cisco, Dell, Lockheed
Martin, Verizon, Uber, Salesforce, Adobe, Qualcomm. This should only include people
speaking directly on behalf of these companies, not people or organizations merely
affiliated with them.

  - "Small Tech": Includes:
      • Technology companies that are not Big Tech
      • Venture capital firms (e.g., a16z, Sequoia)

  - "Civil Society": Includes individuals, nonprofits, academics, media outlets, think
tanks, governments, and all trade associations or councils, regardless of the industry
they represent.
    Also include any large companies that are not primarily technology companies --
such as JPMorgan Chase, Deloitte, Boeing, Pepsi, or others whose core business is not
tech.

  Examples:
  - "OpenAI" → Big Tech
  - "Stability AI" or "Sequoia Capital" → Small Tech
  - "Brookings Institution", "American Chemical Society", or "AI Infrastructure
Alliance" → Civil Society
```

```
  - "Deloitte" or "JPMorgan" → Civil Society

4. fine_tags: One or more labels from the following list:
  "Startup", "Academia", "Nonprofit", "Government", "Individual", "Trade Group",
"Media", "Other".

5. ai_worldview: Return a list of the top two worldviews (primary and secondary) from
this set. Use contextual reasoning to justify your selection.

  ["Accelerationist", "Responsible AI Advocate", "Open AI Innovator", "Safety
Advocate", "Public Interest AI", "National Security Hawk"]

  Worldview Definitions:

  - Accelerationist - Emphasizes speed, innovation, deregulation, and competitiveness
in AI development. Often sees regulation as a barrier to progress.
    Common phrases: "maintain global dominance," "cut red tape," "boost productivity,"
"compete with China"

  - Responsible AI Advocate - Develops AI ethically, with emphasis on fairness,
accountability, inclusivity, and transparency.
    Common phrases: "bias mitigation," "AI fairness," "regulatory frameworks,"
"institutional accountability," "secure model training," "human oversight"

  - Open AI Innovator - Supports open-source AI, open data, and public access to AI
infrastructure and data. Values collaboration, decentralization, and transparency as a
means to democratize AI development.
    Common phrases: "open-source," "democratize," "shared datasets," "open weights,"
"community-led development," "open research," "decentralized development," "model
sharing"

  - Safety Advocate - Prioritizes minimizing long-term and catastrophic risks from
advanced AI. Advocates for alignment, deployment safeguards, and safe systems to
prevent unintended consequences.
   Common phrases: "AI safety," "long-term risks," "existential harm," "AI alignment,"
"containment," "robust safeguards," "safe deployment"

  - Public Interest AI - Seeks AI systems that advance public welfare, inclusion, and
civil liberties. Focuses on accessibility, equitable outcomes, and protecting
individual rights in AI use.
    Common phrases: "rights," "consent," "community," "public good," "data ownership,"
"opt-out," "equity," "creators," "data justice,"" "accessibility," "surveillance
reform," "inclusion"
  - National Security Hawk - Views AI as a strategic asset in global competition and
national defense. Prioritizes protecting critical infrastructure, economic
sovereignty, and leadership in the AI arms race.
    Common phrases: "cyberwarfare," "defense advantage," "military competition,"
"state surveillance," "geopolitical power," "economic sovereignty," "national defense
infrastructure," "global technology race," "export control," "export controls",
"safety act for AI", "defense" , "intelligence agency", "intelligence agencies" ,
"emerging threats", "adversary", "adversarial", "containment," "national
competitiveness," "strategic advantage," "national interests"

  Clarification Heuristics:

  - If the focus is on rights, consent, or creative access, prefer Public Interest AI
over Responsible AI.
  - If the emphasis is on fairness for creators or civilian impact, choose Public
Interest AI rather than Accelerationist.
  - Avoid assigning Accelerationist when the frustration with bureaucracy is not
paired with a strong pro-growth or anti-regulation stance.
  - If the comment is concerned with who benefits from AI and who is left out, prefer
Public Interest AI.
```

```
  - If transparency is discussed in the context of open access, decentralization, or
open-source technology, prefer Open AI Innovator over Responsible AI.
  - If transparency is emphasized alongside fairness, regulation, or accountability,
it aligns more closely with Responsible AI Advocate.
  - Before choosing a worldview, consider whether another label may be more precise.
For example, both Open AI Innovator and Responsible AI Advocate value transparency but
differ in what it is meant to achieve.
  - If the comment emphasizes American AI leadership in the context of international
competition, strategic advantage, or national interest, classify it as National
Security Hawk. This includes framing about a global AI race, export controls,
technological dominance, or economic sovereignty.
  - Do not classify these types of comments as Responsible AI, even if they mention
risk governance or accountability, unless the focus is clearly on fairness or ethical
AI.

6. group_justification: A one-sentence explanation for how you assigned the submitter
to the main group.
7. worldview_justification: A one-sentence rationale or quote from the comment that
supports the AI worldview classification.
8. filter: Is this comment "keep" or "discard"? Keep it if it presents a substantive
opinion on AI policy; discard it if it is unserious, off-topic, or too vague.

Return your final answer ONLY in the following JSON format:
{
 "submitter_name": "...",
 "submitter_type": "...",
 "main_group": "...",
 "fine_tags": [...],
 "ai_worldview": ["...", "..."],
 "group_justification": "...",
 "worldview_justification": "...",
 "filter": "keep"
}
```

Post-Processing and Human Review

After the automated classification stage, we manually corrected a few misclassifications. For example, a comment from a nonprofit organization composed of tech companies was initially labeled as Big Tech and reassigned to Civil Society. Another, from an engineer at Microsoft, was misclassified as Big Tech due to the author's affiliation. However, since the comment is the author's own opinion and not representative of Microsoft's overall stance, we reclassified it to Civil Society.

Final Label Distributions and Visualization

After adjustments, our final dataset included 16 comments from Big Tech, 249 comments from Small Tech, 9,318 comments from Civil Society, and 485 discarded comments which were unserious, off-topic, or too vague. In total, we recorded 18,703 AI worldview classifications across 9,583 comments.

While most submissions included both labels, 463 Civil Society responses lacked a valid secondary classification. These responses do not include detailed policy recommendations but

send a strong message. They express deep skepticism and strong opposition to the current direction of AI development, especially generative AI. A dominant theme is the view that AI relies on the unethical use of copyrighted creative works, framed as theft. Many dismiss AI as overhyped and economically harmful, citing fears of job loss, creative devaluation, environmental impact, and national security risks. Few acknowledge AI's benefits; instead, they call for stronger legal protections, stricter regulation, or even a halt to further development. The tone is broadly distrustful and signals a public demand for more accountability, fairness, and restraint in federal AI policy.

To summarize the worldview patterns across organizational types, we created visualizations using Python's Matplotlib and Seaborn libraries.

The JSON data extracted from the RFI responses can be downloaded from the Harvard Dataverse repository. (Carvao, 2025a)

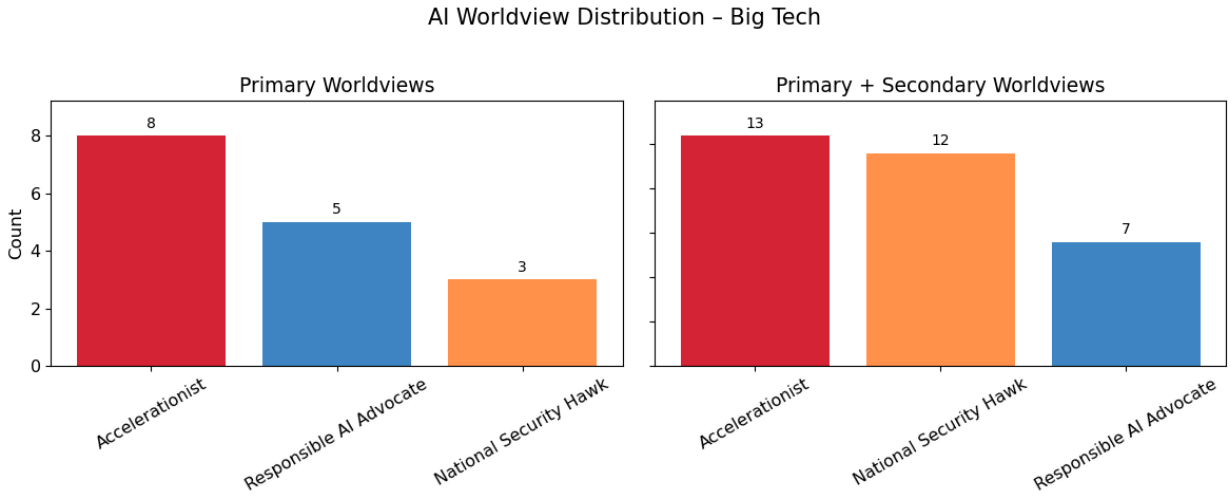

Figure 1 - Big Tech Worldviews

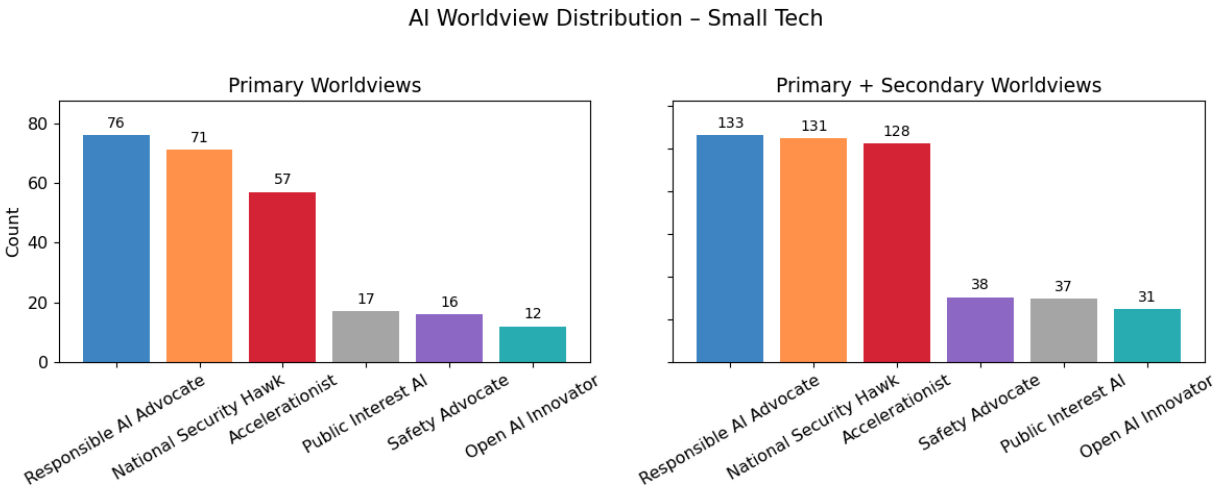

Figure 2 - Small Tech Worldviews

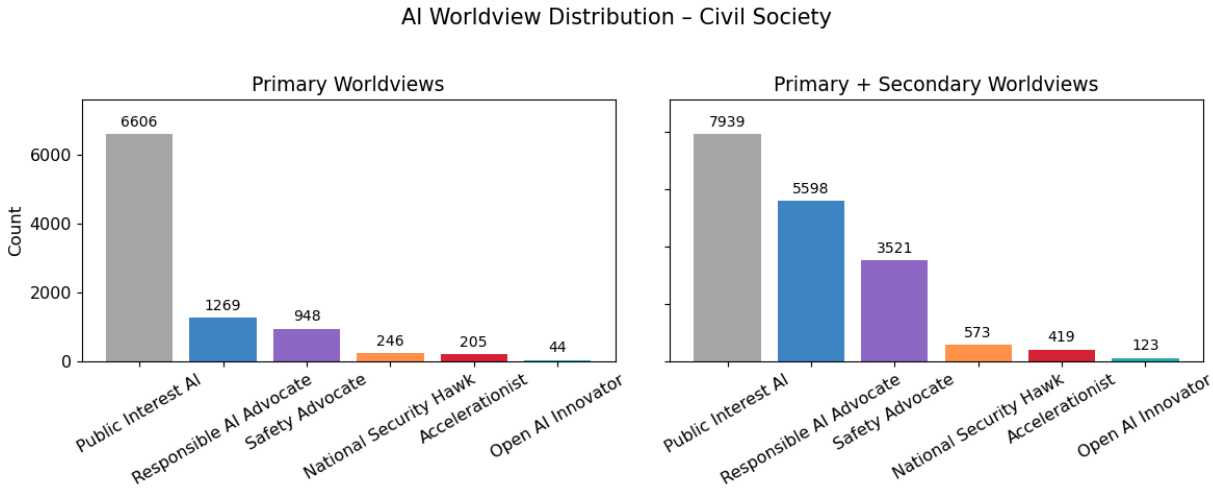

Figure 3 - Civil Society Worldviews

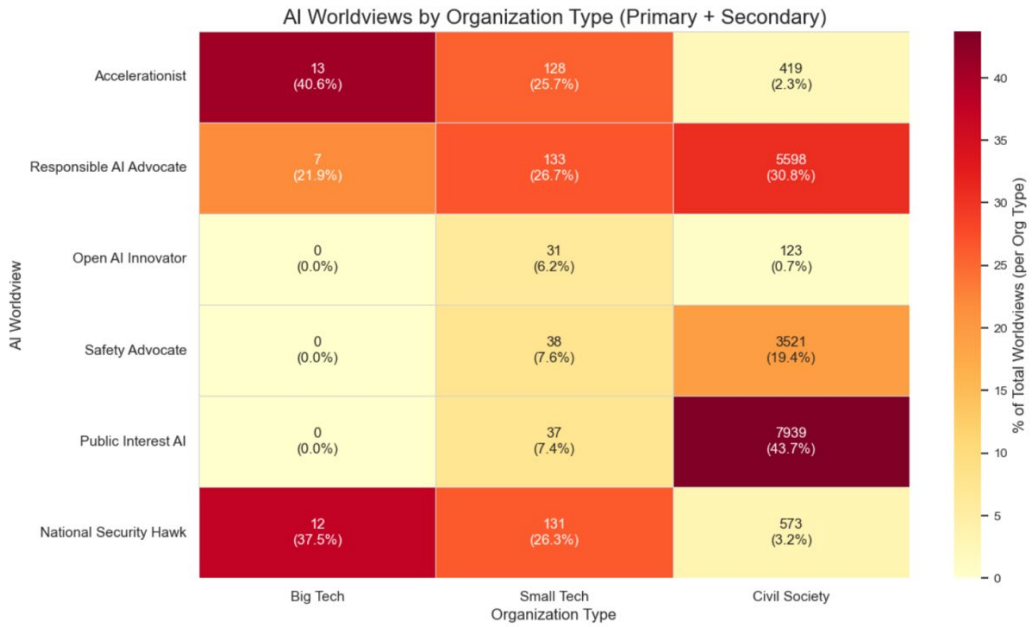

Figure 4 - Worldviews by Organization Type Heatmap

Analysis and Results

Over 10,000 responses present a complex landscape of stakeholder priorities, concerns, and strategic visions. To structure the analysis, we used the six AI worldviews from "Governance at a Crossroads." (Carvao et al., 2025) These archetypes surface the underlying heterogeneity, tensions, and tradeoffs driving AI governance discourse in the nation. All referenced responses are available on the NITRD website (Networking and Information Technology Research and Development, 2025) and appear italicized throughout the text.

Different Worldviews

Accelerationists advocate for rapid AI development with minimal regulatory interference, prioritizing speed-to-market, global competitiveness, and economic growth. Big Tech firms such as OpenAI, Amazon, and Meta exemplify this perspective, with OpenAI urging federal preemption to avoid a *"patchwork of regulations that risks bogging down innovation"* (OpenAI, RFI, p.2, p.6). Meta ties deregulation directly to national growth, asserting that *"[o]pen source AI is the foundation of continued U.S. leadership"* and condemning state-level regulatory complexity as a barrier to progress (Meta RFI, p.1, p.11–12). VC actors like True Ventures, NVCA, and Andreessen Horowitz echo this logic more aggressively, calling algorithmic impact assessments a *"gut punch to startups"* and advocating for a *"light-touch"* regulatory model (True Ventures RFI, p.3–5). Andreessen Horowitz seeks to shield developers from *"onerous regulatory hoops based on speculative fear,"* reinforcing the belief that innovation should not be constrained until concrete harms emerge (Andreessen Horowitz RFI, p.7). The General Catalyst Institute (GCI) blends

deregulatory preferences with infrastructure investment and federal coordination to support startups. They argue that *"a federal, whole-of-government AI policy that sets a national standard and preempts state laws is essential to prevent a fragmented patchwork of state regulations,"* particularly for *"startups with limited resources"* (GCI RFI, p.7). GCI views the current proliferation of local laws as a threat to national coherence and calls upon the federal government to work with the private sector in *"a level of coordination that is without modern precedent"* (GCI RFI, p.5). Workforce development and applied research, increased investment into *"AI Centers of Excellence,"* and partnerships with *"community colleges and trade schools to integrate AI into the existing curriculum"* are additional priorities for GCI (GCI RFI, p.6–7). Unlike the purely deregulatory framing seen in Andreessen Horowitz's response, GCI's accelerationism promotes not only speed but also scale, education, and national resilience. Though the Accelerationist camp faces challenges around ethical risks, safety concerns, and societal impacts due to limited oversight, these accelerationist actors jointly frame AI as a global race that the U.S. cannot afford to lose.

Responsible AI Advocates are concerned with developing ethical, fair, and transparent AI systems, emphasizing procedural accountability and technical documentation. Big Tech firms like Microsoft outline a foundational trust-based paradigm, arguing that *"trust in how technology itself performs"* and *"trust that people and organizations can deploy it successfully"* are prerequisites for broad AI adoption, highlighting both technical and social forms of trust (Microsoft RFI, p.14). IBM similarly champions a governance model rooted in documentation, data lineage, and clarity of roles across the AI supply chain. They encourage using voluntary impact assessments to evaluate high-risk systems instead of mandatory third-party audits through "*transparency governance"* and *"best practices for data provenance"* (IBM RFI, p.7–8). Microsoft and IBM support risk-based, sector-specific regulation grounded in existing legal domestic frameworks, citing the National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF) as a governance touchstone. Intel globalizes this logic, advocating for frameworks "*informed by international standards"* such as those developed by the International Organization for Standardization (ISO) or International Electrotechnical Commission (IEC) (Intel RFI, p.1). Their support for the unified governance is supplemented by a call for hardware-based AI security features, including confidential computing, secure model training, and auditability. Google contributes to this camp by supporting *"model cards and technical reports"* to enhance transparency and usability without forcing trade secret disclosures (Google RFI, p.11). Noting that "*autocorrect features don't pose the same risks (or benefits) as healthcare applications deployed in an emergency room,"* they also align with other Responsible AI Advocates on the stance that AI regulation must be proportional to domain-specific risks (Google RFI, p.7).

Among industry associations, the Business Software Alliance (BSA) backs voluntary accountability mechanisms and alignment with international standards bodies by warning against duplicative, *"confusing and inefficient"* compliance environments, urging for *"one, harmonized,*

flexible, and workable standard for AI governance" to support innovation and public trust (BSA RFI, p.1).

Together, Responsible AI Advocates propose a framework of adaptive, voluntary, and technically grounded governance, sharing a belief that AI regulation should promote transparency and accountability through best practices rather than strict oversight. However, these actors are challenged to balance their aversion to rigid mandates with the growing demand for enforceable redress and public assurances of fairness. Their key strategies of risk scoping, documentation, and self-testing demonstrate a commitment to responsibility, but they often struggle to address systemic bias concerns raised by Public Interest AI actors.

To enable AI transparency, equity, and competitiveness, Open AI Innovators commit to open-source model development, data accessibility, and cross-sector collaboration. Actors in this group promote open ecosystems not simply as technical alternatives to closed models, but as strategic, economic, and governance imperatives to drive innovation. IBM makes a particularly strong case, arguing that *"open model weights will be key to enabling new technical breakthroughs that increase safety"*, citing *"self-destructing models"* as an example of open innovation (IBM RFI, p.3). Open models, in their view, "*encourage greater involvement and scrutiny from a larger community of stakeholders, increasing the likelihood that bias and vulnerabilities are identified and patched"* (IBM RFI, p.1). By stating that overly restrictive regulation of open-source systems *"will only slow our ability to win this technological race"* and hinder U.S. soft power in global AI leadership, they invoke that national competitiveness is made possible through such open systems (IBM RFI, p.5). Meta likewise places openness at the core of its economic and national security strategy by positioning open-source models in federal procurement, especially in sensitive domains like national security and logistics, as foundational for the *"secure processing of sensitive data"* (Meta RFI, p.4). They also oppose export controls on open model weights, claiming such measures would *"only serve to weaken U.S. interests"* while advantaging geopolitical rivals like China (Meta RFI, p.3).

Intel and the GDA focus on openness in the context of data infrastructure and international regulatory harmonization. Intel endorses *"open ecosystems, including open-source software,"* to lower barriers to entry for AI adoption. The GDA emphasizes the necessity of "*cross-border access to, and transfers of, data"* to sustain innovation, calling for *"interoperable data transfer mechanisms"* against data localization requirements that may fragment the global AI economy (Intel RFI, p.3; GDA RFI, p.2). The U.S. Chamber of Commerce recommends investment in shared infrastructure like the National AI Research Resource (NAIRR) to support *"the open-source ecosystem"* (Chamber of Commerce RFI, p.8). Their advocacy combines pro-openness messaging with intellectual property protections and global standards alignment, representing a balance between innovation acceleration and strategic control.

Though these actors vary in the degree to which they view openness as a geopolitical instrument, a technical safety mechanism, or a market equalizer, they are ultimately united by a belief in

openness as a force multiplier to drive innovation, reduce redundancy, and expand participation. Tensions remain around the potential for misuse, intellectual property vulnerabilities, and the geopolitical consequences of unbounded diffusion of newly developed AI models. These debates illustrate that within the open paradigm, optimism for technical innovation often coexists with caution.

Safety Advocates constitute a smaller but tightly knit coalition in the AI governance space committed to long-term risk mitigation, alignment, and the responsible deployment of increasingly capable frontier models. Unlike the open innovation or market-first logic embraced by open-source proponents and accelerationists, Safety Advocates prioritize enforceable standards, infrastructure hardening, and pre-deployment evaluation as prerequisites to new model development. Anthropic is the most vocal member of this camp, advocating for mandatory model evaluations for national security-relevant capabilities and proposing that the federal government *"partner with industry leaders to substantially enhance security protocols at frontier AI laboratories to prevent adversarial misuse and abuse*" (Anthropic RFI, p.5–6). Unlike traditional enterprise big technology suppliers, Anthropic's protectionist stance places safety and security ahead of deployment velocity, warning that even the theft of a single frontier model could *"collapse the export control regime"* (Anthropic RFI, p.6). They emphasize that *"if there is evidence that AI systems pose critical national security risks, then developers like Anthropic should be required to test their systems for these risks"* (Anthropic RFI, p.3). By suggesting that the federal government be equipped with rigorous evaluative capabilities and that the AI Safety Institute within the Department of Commerce be preserved, they seek to build out Memoranda of Understanding (MOUs) with AI firms developing *"comprehensive national security evaluations for powerful AI models"* in collaboration with NIST, the Intelligence Community, and the Department of Defense (Anthropic RFI, p.4).

Intel adds an infrastructure-focused perspective, advocating for the deployment of hardware-based security measures like *"Trusted Execution Environments (TEE)"* and machine-readable provenance tools for supply chain transparency (Intel RFI, p.5). They further endorse federated learning to enhance privacy, stating that such will ensure that *"data and computations remain secure and in the owner's custody, even in untrusted environments,"* and should be thus *"uniformly available across different processor types to maintain consistent protection"* (Intel RFI, p.5).

Anthropic and Intel argue that safety must be embedded in AI governance and infrastructure. Their focus extends beyond responsible principles to enforceable protocols, interagency coordination, and secure-by-design systems. A central challenge is defining enforceable standards that keep pace with rapid advances and secure industry consensus. Safety Advocates navigate national security and innovation trade-offs, especially where model testing and export controls may hinder agility. Still, their position offers a critical counterweight to accelerationist extremes, grounding the debate in precaution.

The trope of **Public Interest AI** presents a fundamentally different paradigm than most commercial actors. Rather than focusing on innovation, speed, or competitiveness, stakeholders including the Leadership Conference on Civil and Human Rights, the Electronic Frontier Foundation (EFF), and Harvard Kennedy School's (HKS) Allen Lab for Democracy Renovation, center their analysis on equity, democratic integrity, accountability, and redress. These actors argue that without significant structural safeguards, AI will exacerbate existing injustices and undermine humanity. The Leadership Conference details a comprehensive critique of how AI deployment has *"led to the loss of economic opportunities, higher costs or denial of loans and credit, wrongful arrests, adverse impact on individuals' employment or ability to get a job, lower quality healthcare, and barriers to housing,"* for marginalized communities (Leadership Conference RFI, p.3). Asserting that *"[n]one of the desired outcomes in those areas are achievable if AI systems are not trustworthy and if they are not safe, effective, and fair,"* they proposed that *"human-centered design (HCD)"* systems should be made with legally enforceable bias audits and impact assessments to ensure transparency and accountability (Leadership Conference RFI, p.2). The U.S. Technology Policy Committee's recommendations center on redress as a necessary mechanism to secure justice for those adversely affected, calling for clear documentation of data sources, algorithms, and decision-making processes so that *"users and oversight bodies can trace how outcomes are reached"* (USTPC RFI, p.2).

Meanwhile, the EFF focuses on government accountability and transparency, warning that unregulated procurement and deployment of AI in law enforcement and surveillance contexts risks systemic abuse. The EFF views *"public notice-and-comment practice[s]"* as essential to *"weed out bogus products and identify applications where certain types of tools, such as AI, are inappropriate"* (EFF RFI, p.2). They uniquely prioritize open governance, rather than technical improvement, as the primary criterion of safety.

The HKS Allen Lab for Democracy Renovation explicitly warns that without critical intervention, "*AI's ability to manufacture consensus and manipulate voting behavior will pose an existential challenge to free societies and national security"* (GETTING-Plurality RFI, p.3). Their framework is premised on the democratic principle that AI should support, not manipulate, public discourse, describing scenarios in which AI could *"fabricate political scandals, alter public records, and misrepresent candidates,"* if regulatory responses remain reactive (GETTING-Plurality RFI, p.3).

To Public Interest AI Actors, safety is intrinsically linked to distributive justice, transparency, and the defense of democratic norms. These groups demand structural reforms grounded in the rights of individuals, particularly those historically excluded from technological design and decision-making. Their challenges include limited funding and the dominance of commercial interests in AI governance, which continues to marginalize public-interest perspectives.

National Security Hawks treat AI as a strategic asset critical to U.S. geopolitical power, economic sovereignty, and defense infrastructure. These actors frame AI governance not only as a matter of domestic competitiveness but as a lever in a rising global technology race.

Microsoft and Amazon adopt a strategic infrastructure approach to national defense, prioritizing AI integration across classified environments. Microsoft calls for the full *"adoption and use of domestic AI models by U.S. defense and intelligence agencies for mission enablement and efficiency"* and underscores the urgency for preparing for *"emerging threats"* posed by adversaries such as China, Iran, North Korea, and Russia (Microsoft RFI, p.8, p.10). They highlight that their investment in *"classified cloud infrastructure"* allows agencies to access *"the most advanced AI models at scale"* (Microsoft RFI, p.8). Amazon similarly promotes its AWS Top Secret Cloud and Bedrock's integration into national security systems, arguing that *"[m]aintaining AI leadership in national security and defense requires increased and sustained AI infrastructure investment in government-specific and classified environments"* (Amazon RFI, p.5).

Anthropic represents a highly proactive and protectionist stance, calling for strict export controls, government-mandated model classification, and public-private partnerships to *"prevent adversarial misuse and abuse of powerful AI technologies"* (Anthropic RFI, p.6). To mitigate these threats, they propose government-to-government treaties that compel nations that are *"high-risk for chip smuggling"* to align their export controls with the U.S., creating a system of diplomatic technical containment (Anthropic RFI, p.5).

This coalition of National Hawks splinters on the question of openness. Meta frames openness as a geopolitical strength rather than a liability claiming it "*reduces the barriers to entry to AI"* and *"drives down prices so that AI can be deployed not only by the few technology companies that have the computing infrastructure to pretrain large models, but also by the broader community of developers"* (Meta RFI, p.6). They urge the federal government to support the export of open models to extend American technical standards globally. In contrast, Anthropic and Amazon view openness as a direct national security risk. Both advocate for controlled diffusion of model weights and capabilities, with Anthropic calling for formal restrictions on powerful models outside aligned nations and Amazon warning that AI deployment must meet the highest levels of security. The U.S. Chamber of Commerce attempts to broker a middle ground, supporting targeted export controls to *"limit adversary access to critical technologies,"* while cautioning that *"excessively burdensome requirements"* could create *"competitive disadvantages for U.S. companies"* (Chamber of Commerce RFI, p.11).

The group of national security hawks is not a monolith, but what unites them is a consensus that AI is not just a commercial tool, but a key to geopolitical competition.

Converging Concerns and Diverging Values

Among these archetypes, a few themes emerge. First, **a tension between acceleration and safety.** Accelerationists like OpenAI downplay safety in favor of rapid deployment of new systems, while safety advocates like Anthropic argue for binding regulation and red teaming of high-risk models. Traditional enterprise providers such as IBM, Microsoft, and Intel straddle both sides, advocating for rapid infrastructure investment while simultaneously championing trust, transparency. These firms urge regulatory restraint to preserve innovation momentum.

These tensions reflect deeper structural divides. Big technology firms and industry coalitions are largely focused on preserving global competitiveness, workforce productivity, and infrastructure deployment. Corporations frame safety as a technical property that can be standardized through public-private partnerships, incentives, and proper procurement. Civil society, by contrast, sees safety as intrinsically bound to equity and accountability towards a broader audience beyond simply manufacturers and governments. Their recommendations often include community oversight, legal liability, and transparent grievance processes. While the industry asks if AI systems can be built safely, civil society asks for whom these systems would be considered safe, and how this safety is defined.

Second, the debate between **federal preemption and state regulation** is a defining fault line between industry and civil society. Many industry actors, particularly Accelerationists like True Ventures, Google, and Meta, warn against a patchwork of state laws, advocating for a national framework to enable compliance certainty and reduce friction. They argue for federal preemption to avoid regulatory havoc from state-by-state compliance, while civil society organizations tend to welcome local experimentation and citizen-led policymaking. Google, for instance, notes that the administration should ensure that *"the U.S. avoids a fragmented regulatory environment that would slow the development of AI, including by supporting federal preemption of state-level laws that would affect frontier AI models"* (Google RFI, p.5). Differently, the EFF notes that *"state legislatures around the country [..] have begun to grapple with questions of fairness and legal compliance"* (EFF RFI, p.3). These perspectives expose the tension between national scalability and localized responsiveness.

Third, **infrastructure and capacity building** emerge as a nearly universal concept amongst Big Tech firms such as Amazon, Microsoft, and Anthropic, which each highlight the need for massive public-private investment in energy, computing data infrastructure, and AI workforce development. Microsoft, for instance, has committed to *"investing over $50 billion in America's AI infrastructure in 2025"* (Microsoft RFI, p.1), while Anthropic projects that *"by 2027, training a single frontier AI model will require networked computing clusters drawing approximately five gigawatts of power"* (Anthropic RFI, p.6).

Fourth, **accountability and redress** are predominantly championed by civil society organizations that believe that AI should work for the sake of everyone in society, calling for safeguards against bias and redress mechanisms. The Leadership Conference, for example, demands regular bias audits and clear documentation of model decision-making (Leadership Conference RFI, p.3), while EFF emphasizes the need for legally mandated transparency in AI procurement (EFF RFI, p.2). Industry respondents rarely mention algorithmic discrimination or equity, signaling a stark separation in belief regarding these camps.

Finally, persistent **open-source tensions** exist within Big Tech. IBM and Meta argue that open-source models foster innovation and safety via community scrutiny (IBM RFI, p.3; Meta RFI, p.1), while Anthropic *"strongly recommend[s] the administration strengthen export controls on*

computational resources and implement appropriate export restrictions on certain model weights" (Anthropic RFI, p.4) and Amazon view openness as a vulnerability urging the government to consider *"how AI can support efforts to enhance cyber defense."* (Amazon RFI, p.5). OpenAI highlights that *"the question of whether AI should be open or closed source is a false choice—we need both, and they can work in a complementary way that encourages the building of AI on American rails."* (OpenAI RFI, p.7-8)

Patterns of Influence

While the RFI responses contain a great diversity of opinion, the six structured archetypes offer a lens through which to understand the competing interests shaping U.S. AI policy. The AI Action Plan will need to balance innovation with safety, centralization with state regulation, and growth with public transparency.

The industry does not speak with one voice. Positions vary by business model, size, and perceived strategic advantage. While virtually all industry actors agree that the U.S. should maintain a leadership position in AI development, they express divergent views on how to get there, how much regulation is appropriate, and what types of innovation should be prioritized.

Among **large technology firms**, OpenAI, Meta, and Amazon exhibit an accelerationist orientation. OpenAI's submission, for instance, opens with an admonition against *"layers of laws and bureaucracy"* that could inhibit rapid deployment and warns that a patchwork of regulations risks slowing down innovation, *"undermining America's leadership position"* (OpenAI RFI, p.2, p.6). Meta echoes this concern and calls for federal preemption to avoid the *"unworkable regulatory environment"* created by divergent state laws (Meta RFI, pp.11-12). Amazon, likewise, focuses on scaling energy infrastructure and removing permitting barriers, asserting that preserving U.S. leadership in generative AI and emerging technology is dependent upon efforts to prioritize *"new investments in transmission and generation infrastructure"* in short to *"scale with the demands AI workloads require"* (Amazon RFI, p.3).

T**raditional enterprise technology firms** like Microsoft and IBM walk a more nuanced line, still emphasizing innovation, while framing trust and accountability as necessary enablers of the adoption of AI. Microsoft explicitly notes that *"to create the conditions for innovation and AI adoption at scale,"* one needs durable public support and states that "*[p]roperly resourcing NIST is vital to ensuring the United States continues to lead in developing the measurement backed science and standards that will facilitate [..] adoption."* (Microsoft RFI, pp.13-14). IBM similarly endorses *"clear differentiations between obligations for developers and deployers"* to encourage AI innovation and protect consumers by placing controls on *"specific uses of AI that pose the greatest risk to people and their wellbeing"* (IBM RFI, p.7). These Responsible AI Advocate actors seek a balance, promoting innovation without sacrificing accountability.

The **Small Tech category** reveals a more diverse set of positions. Workday and ServiceNow align most closely with Responsible AI Advocates, particularly through their championing of the NIST

AI RMF. Workday emphasizes that their implementation *"was NIST's first-ever case study highlighting industry adoption"* (Workday RFI, p.3). ServiceNow echoes this support, calling for *"voluntary—rather than compulsory—standards"* that include "*testing, evaluation, security, and documentation"* (ServiceNow RFI, p.1). Cloudflare offers a strong endorsement of open-source, arguing that "*to maintain U.S. leadership in AI, the government should permit the free global transfer of open-source models"* and frames openness as a geopolitical asset (Cloudflare RFI, p.5). Snowflake similarly embraces Open AI Innovator values, demanding that "at least 50 high-impact data sets that will make an outsized impact on AI Development" should be identified and made available to researchers and innovators," promoting an *"international data sharing arrangement to hasten the development of AI"* (Snowflake RFI, p.6). CrowdStrike presents a national security-driven perspective warning that AI has already enabled *"unsophisticated threat actors to achieve nation-state level cyber capabilities"* and that understanding adversary use is *"critical to maintaining the cybersecurity advantage"* (CrowdStrike RFI, p.3).

Small Tech's accelerationist energy is largely derived from the influence of VCs. True Ventures, NVCA, and Andreessen Horowitz adopt an aggressive deregulatory stance, opposing state-level algorithmic impact assessments and federal safety proposals. True Ventures adamantly argues against *"onerous liability regimes designed to hold developers financially responsible for downstream harms"* because of their detriment to startups unable to afford significant compliance budgets (True Ventures RFI, p.3). They warn of a *"50-state AI regulatory morass"* that could crush bottom-up innovation (True Ventures RFI, p.5). NVCA calls for massive public investment in R&D to support startup-led innovation, linking national competitiveness directly to startup liquidity and acquisition pathways. Notably, they call for the continuation of bipartisan *"federal grants and programs such as Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) [to] provide non-dilutive capital to help startups experiment with novel AI models and applications without giving up equity."* (NVCA RFI, p.4). Andreessen Horowitz urges that the government *"prohibit the harms and punish the bad actors"* but avoid speculative guardrails that *"cripple competition"* before evidence of actual misuse arises (Andreessen Horowitz, p.8). Beyond liability, Andreessen Horowitz expands the conversation about federal preemption, stating that *"[g]iven that the AI model market is national in scope and directly implicates interstate commerce, the federal government should take the lead in its promotion and regulation to provide uniformity and certainty for US model developers."* (Andreessen Horowitz RFI, p.5)

Industry associations and Trade Groups such as the U.S. Chamber of Commerce, Business Roundtable, and BSA simultaneously champion federal preemption, global interoperability, and infrastructure investment and advocate for voluntary, risk-based frameworks over binding mandates. The BSA warns that inconsistent rules from *"each regulatory agency and all 50 states"* would *"confuse and inhibit"* both developers and adopters, pushing for a single harmonized standard (BSA RFI, p.1). The U.S. Chamber of Commerce highlights the role of AI in empowering small and medium-sized enterprises, citing that *"98% of small businesses are using some form of*

an AI tool" and that these businesses experience higher rates of growth when AI is adopted (U.S. Chamber of Commerce RFI, p.1-2). They tend to reflect what established technology companies think, instead of the strong accelerationist stance from VCs.

A different pattern of responsible AI advocacy emerges from **other non-technology companies that we grouped with Civil Society**. Johnson & Johnson emphasizes the need for the promulgation of "*voluntary—rather than compulsory–standards"* to *"accelerate AI innovation, adoption, and collaboration*" in healthcare, framing AI development around clinical validity and results-based, "*meaningful outcomes"* (JnJ RFI, p.3). Similarly, PepsiCo highlights workforce development, cybersecurity safeguards, and sustainable deployment, noting that *"although some employers may view AI and automation as immediate solutions, deployment must be done responsibly"* with the *"dual goals of improving worker autonomy and adapting to new technologies"* (PepsiCo RFI, p.2). JPMC commits to responsible deployment, noting that *"multistakeholder standards development is essential for fostering innovation and ensuring that AI technologies are used responsibly"* (JPMC RFI, p.8). They acknowledge the necessity to manage the risks of AI, noting that they have a *"thorough, well-established risk and control framework"* to *"serve clients and communities effectively and responsibly."* (JPMC RFI, p. 2).

Academia, advocacy associations, and think tanks such as the HKS Allen Lab, the Leadership Conference on Civil and Human Rights, and the ACM U.S. Technology Policy Committee embody the Public Interest AI archetype, framing AI governance around civic resilience, democratic equity, and structural harm prevention. The Electronic Frontier Foundation explicitly warns that AI "reproduces historic bias by design" and calls for civil rights audits, transparency, and redress mechanisms as preconditions to responsible use (EFF RFI, p.2). The HKS Allen Lab similarly warns of AI's ability to *"manufacture consensus,*" calling for investment in public infrastructure and civic education (GETTING-Plurality RFI, p.3). These perspectives mark a striking divergence from the commercial sector, prioritizing accountability, public deliberation, and institutional safeguards over innovation speed.

Together, these patterns reveal that while the industry aligns to promote national AI leadership, it diverges on how this should be accomplished. Frontier labs emphasize speed and competitiveness, traditional large IT infrastructure firms call for trust-building and safety mechanisms, and venture capital firms seek deregulatory protections for startups. Other key industry associations negotiate between the pressing imperatives for growth and regulatory pragmatism. Civil society actors, especially those affiliated with academia and advocacy, frame governance through the perspective of rights, transparency, and social equity. The result of these differing beliefs is a complex and often contradictory policy environment.

The Dividing Lines

While Big Tech and VC firms seek to preserve U.S. leadership in AI and reject restrictive regulation, their approaches, constituencies, and governance proposals diverge in meaningful

ways, particularly when compared to the priorities of civil society. On one level, Big Tech and VC actors converge on the importance of federal preemption. Google seeks to avoid *"a fragmented regulatory environment that would slow the development of AI,"* while Meta similarly highlights the dangers of "*an unworkable regulatory environment"* (Google RFI, p.5; Meta RFI, p.11–12). VCs echo this stance. True Ventures criticizes state-level impact assessments, contending they would "*significantly restrain American competitiveness"* and delay startup innovation. GCI likewise calls for a *"whole-of-government"* approach to avoid state-by-state compliance burdens (True Ventures RFI, p.3–5; GCI RFI, p.7). Despite this policy convergence, the two are driven by distinct motivations. Big Tech seeks scale and operational certainty. VCs, in contrast, aim to lower entry barriers for the startup ecosystem.

These groups split sharply in their governance philosophies. Firms like Microsoft and Amazon endorse trust-based, voluntary frameworks such as the NIST AI RMF. Microsoft stresses the importance of building public trust and recommends investment in shared infrastructure like the NAIRR. Amazon similarly calls for *"appropriate risk-based governance practices"* developed with federal agencies (Amazon RFI, p.6). By contrast, Andreessen Horowitz frames proactive compliance frameworks as speculative, warning that "*complex compliance regimes based on the math that an engineer uses"* will deter innovation by *"Little Tech"* but still *"not change whether a criminal is held liable when they use AI to commit fraud"* (Andreessen Horowitz RFI, p.7–8).

In the National Security domain, Big Tech and VCs diverge again. Microsoft and Amazon support the integration of AI into defense infrastructure, citing the need for classified cloud environments and public-private collaboration for mission-critical applications. Anthropic adds a protectionist angle, calling for model evaluations, export controls, and international agreements to prevent adversarial misuse. Uniquely, VCs like True Ventures propose a *"SAFETY Act for AI"* to limit liability, alongside the creation of an *"AI Information Sharing and Analysis Center (ISAC)"* for collective cybersecurity threat intelligence (True Ventures RFI, pp.6-7). GCI looks at national security through a lens of infrastructure modernization, stating that: *"AI supply chain security and protections against corporate espionage are critical"* (GCI RFI, p.6).

In sharp contrast, civil society groups reject national security as a dominant paradigm for AI governance. The EFF critiques the opaque procurement process and urges for *"public notice-and-comment"* procedures to safeguard civil liberties (EFF RFI, p.2). The HKS Allen Lab warns that unchecked AI deployment could represent a threat not only to U.S. sovereignty but to democracy itself. These groups position safety as a question of procedural fairness, accountability, and public transparency, not as a matter of infrastructure resilience or adversary deterrence.

In summary, the 2025 RFI responses reveal a deeply stratified and fragmented AI governance landscape where shared ideas for U.S. global leadership co-exist with sharp ideological, strategic, and operational divides. The six archetypes (Accelerationists, Responsible AI Advocates, Open AI Innovators, Safety Advocates, Public Interest AI, and National Security Hawks) offer a crucial

lens for understanding these tensions, clarifying not only where the disagreements lie but also the diverse constituencies that these positions represent. Civil society offers a distinctly justice-oriented vision of AI governance. These democratic priorities are largely absent from the submissions of corporations, pointing to deeper disagreements not simply about regulation, but about the very purpose of AI governance. As the administration seeks to construct the U.S. National AI Action Plan, it must resist the urge to impose a singular industry-driven narrative and instead strive to balance innovation with safety, global competitiveness with domestic equity, and pursue AI leadership without compromising the democratic commitments to public accountability.

Discussion

What can be learned from the responses, and what main themes emerge?

By employing a combination of qualitative review and ML-based classification, five key patterns emerged: (1) the tension between innovation acceleration and safety concerns; (2) the debate over federal preemption versus state-level experimentation; (3) broad consensus on infrastructure and capacity-building; (4) divergent views on accountability and redress mechanisms; and (5) differing philosophies surrounding open-source AI development.

Industry responses predominantly emphasized national competitiveness, infrastructure, and regulatory flexibility. Civil society focused on issues of equity, accountability, and democratic oversight. The sheer volume and granularity of the submissions speak to the complexity of the policy landscape: one that demands a governance model capable of integrating plural and sometimes conflicting values.

What is the industry advocating for, and does it speak with one voice?

To assess industry cohesion, we examined responses from Big and Small Tech, industry associations, and VC stakeholders. While there is agreement on the need to promote U.S. global leadership and reduce regulatory fragmentation, substantial divergence exists in preferred governance approaches.

Big Tech largely supports voluntary, risk-based frameworks, infrastructure investment, and public-private partnerships to enable responsible AI deployment. Yet these firms are far from uniform. IBM and Microsoft champion trust, documentation, and transparency, while others like Meta, Amazon, and OpenAI adopt a more accelerationist stance, warning that excessive regulation could threaten innovation and competitiveness.

Small Tech reveals a heterogeneous policy orientation, with firms such as Workday and ServiceNow aligning closely with Responsible AI principles, explicitly endorsing voluntary standards and risk-proportional governance structures. Others like Cloudflare and Snowflake spotlight the value of fostering an open and decentralized AI ecosystem, arguing that open-source models and international data sharing agreements are catalysts for innovation. Crowdstrike, meanwhile, centers on national security, cybersecurity resilience, and adversarial preparedness.

VC firms represent the most deregulatory faction opposing state-level safety mandates and resisting federal compliance regimes. Their submissions frame such measures as burdensome to

early-stage innovation and favor uniform federal preemption to avoid a regulatory morass amongst the states.

Industry associations attempt to bridge this gap and endorse harmonized, voluntary standards while encouraging investment and innovation without strict oversight. The rhetoric from these actors reflects the pragmatic concerns of both large-scale technology firms and VC ecosystem builders, aiming to unify industry voices without alienating varying stakeholders.

This fragmentation underscores that the industry is not a monolith, but a coalition of actors with differing incentives, constituencies, and visions for AI governance.

Is Big Tech's position aligned with the VC community and Small Tech?

Big Tech's position is best characterized as pragmatically deregulatory. Traditional enterprise firms like Microsoft, IBM, and Intel support responsible AI development through voluntary measures, transparency standards, and alignment with international frameworks. Trust-based frameworks and proportional, sector-specific risk governance are viewed as fostering public confidence and innovation by these firms. Conversely, frontier labs and platform companies, such as OpenAI, Amazon, and Meta, adopt an accelerationist approach, favoring preemption and rapid deployment.

Small Tech, which includes VC firms and other non-Big Tech companies, is a more varied cohort. Workday and ServiceNow mirror Big Tech's responsible AI posture. Crowdstrike, by contrast, foregrounds Big Tech's national security concerns by similarly emphasizing adversarial threat preparedness and AI cyber defense. VCs take a markedly deregulatory approach and reject state-level and proactive federal governance frameworks, warning that overly burdensome compliance regimes threaten startup viability.

Thus, while both Big and Small Tech oppose regulatory fragmentation and call for federal preemption, they differ significantly in their governance philosophies. Big Tech is more open to public safety collaboration and trust-building, whereas VC actors prefer minimal government intervention.

Are the industry positions aligned with patterns previously identified?

Yes, the taxonomy developed in "Governance at a Crossroads" proved effective for categorizing the RFI responses:

- **Accelerationists** (OpenAI, Meta, True Ventures) push for speed, scale, and regulatory restraint.
- **Responsible AI Advocates** (IBM, Intel, Microsoft) emphasize trust, transparency, and voluntary oversight.
- **Open AI Innovators** (Meta, IBM, U.S. Chamber) view openness as a key driver of safety and competitiveness.
- **Safety Advocates** (Anthropic, Intel) prioritize security protocols, enforceable standards, and risk mitigation, often at the expense of rapid deployment.
- **National Security Hawks** (Amazon, Anthropic, Microsoft) frame AI as a strategic asset and support defense integration and export controls.

Few industry actors align with the **Public Interest AI** framework, underscoring a values gap between industry and civil society on equity, bias, and democratic accountability.

What positions have been expressed by civil society and universities?

Civil society organizations and academic institutions offer a fundamentally different lens to AI governance. Their approach is rooted in rights, justice, and democratic safeguards with a focus on:

- Structural inequities amplified by AI (e.g., discrimination, surveillance, exclusion).
- The need for legally enforceable redress and bias audits.
- Democratic accountability in AI procurement and deployment.
- The dangers of overreliance on industry self-regulation.

Groups like the Leadership Conference and the EFF critique technical governance alone as insufficient, calling for inclusive, community-oriented systemic safeguards. Academic centers such as the Association of University Presses (AUP), Association of Public and Land-Grant Universities (APLU), and the Association of American Universities (AAU) acknowledge AI's potential to enhance U.S. competitiveness, while urging for stronger academic-industry collaboration. The HKS Allen Lab highlights how AI may distort democratic processes, calling for civic infrastructure that can withstand the manipulative use of generative models.

These responses advance a paradigm of just AI, recognizing that technical excellence alone cannot ensure equitable outcomes. Their advocacy centers on civil rights-based governance structures that are power-sensitive and participatory in design.

Does the industry align with civil society?

Alignment between industry and civil society remains limited and mostly rhetorical. **Industry actors** define safety in operational or technical terms**,** favoring voluntary compliance, secure model development, and interoperable standards. Many view risk as something to be managed through incentive structures. **Civil society**, by contrast, defines safety in ethical, social, and democratic dimensions, emphasizing structural safeguards like bias audits, public oversight, and legal redress mechanisms.

This divide is further reflected in regulatory preferences. Industry consistently advocates for federal preemption to avoid a patchwork of state laws and ensure compliance uniformity. Civil society welcomes local experimentation and citizen-driven participatory governance. Whereas industry frames regulation as a barrier to competitiveness, civil society sees its absence as a threat to justice.

Several non-technology companies classified under Civil Society take a hybrid approach that blends commercial interests with public interest values, showing greater sensitivity to ethical governance principles. Johnson & Johnson stresses explainability, ethical use, and outcome-driven regulation: *"We believe ethical considerations should be at the forefront of how we are applying AI models, and that regulation has a key role to promote the responsible and transparent usage of AI"* (J & J RFI, p.3). JPMC advocates for global regulatory harmonization and voluntary frameworks like the NIST AI RMF to enhance U.S. leadership. PepsiCo integrates responsible governance with public interest approaches, warning against the *"negative consequences"* of

"hasty implementation", and advocating for equitable deployment, workforce integration, and sustainable agriculture partnerships (PepsiCo RFI, p.2).

These actors reveal that the boundary between industry and civil society may be more porous than assumed. Ultimately, the debate reflects two competing paradigms of governance: one that privileges speed, standardization, and voluntary mechanisms; the other prioritizes inclusion, rights-based safeguards, and legal accountability. Bridging this divide will require governance models that integrate both perspectives to ensure AI is not only effective and scalable but also equitable, just, and publicly trusted.

Conclusion

The authors were not surprised to see a divergence of opinion between industry and civil society. There is an opportunity and a need, however, for these perspectives to converge. Artificial Intelligence is likely to be the next general-purpose technology with a profound and economy-wide impact on productivity and growth, and a transformative effect on society. Civil society must reflect on the economic opportunity at hand and the urgency to act swiftly. At the same time, industry must acknowledge the sociotechnical impacts of its creations. The government cannot abdicate its responsibility to reconcile these perspectives. Blind or ideological advocacy for one side or the other is not only inadequate, but it can be harmful in the long run.

The opportunity ahead of us is clear. American society needs to decide how to approach this moment. Regardless of the nature of the soon-to-be-announced AI Action Plan, a decision needs to be made on whether to consolidate a laissez-faire approach to technology (Lima-Strong & Hendrix, 2025) or to establish safeguards to protect against harms. This is a normative judgement that belongs in Congress. Today, the tech industry in the US is subject to antitrust and civil litigation. While effective means to seek harm redress, both methods presuppose the existence of the harm. What if there were types of harms that, as a society, we would prefer to prevent and avoid? This is where sectoral regulation can help. (Carvao, 2025b)

Today's claims for light regulation (Brody, 2025) are a disguised way to promote no regulation and a forceful attempt to drive an ideological opposition to what is seen as excessive and premature governmental intrusion. (Adam Thierer et al., 2025) The debate between no regulation and too much regulation is a false choice. (Carvao, 2024) Alternatives such as AI sandboxes (The Datasphere Initiative, 2025), dynamic governance models (Carvao et al., 2025), or multi-stakeholder regulatory organizations (Jerry McNerney, 2025) are proposals currently on the table.

The development of the U.S. AI Action Plan presents a timely opportunity to align policy with a wide range of stakeholder priorities. The submissions reviewed here illustrate both common ground and significant divergence, particularly between industry and civil society. Understanding these perspectives—and the strategic interests behind them—will be essential for designing a governance framework that is adaptive, credible, and capable of supporting long-term innovation and public trust.

References

Adam Thierer, Greg Lukianoff, & Dean W. Ball. (2025, April 3). How state AI regulations threaten innovation, free speech, and knowledge creation. *R Street Institute*. https://www.rstreet.org/?post_type=commentary&p=85535

Brody, B. (2025, May 8). Cruz to push 'light touch' AI regulation. *Punchbowl News*. https://punchbowl.news/article/tech/cruz-to-push-ai-regulation/

Carvao, P. (2024). The dual imperative: Innovation and regulation in the AI era. *International Journal of Technology Policy and Law*, *3*(3), 236–251. https://doi.org/10.1504/IJTPL.2024.142861

Carvao, P. (2025a). *JSON data extracted from the public comments to the AI Action Plan RFI (Jun 2)* (Version V1) [Dataset]. Harvard Dataverse. https://doi.org/10.7910/DVN/JS8BQV

Carvao, P. (2025b, March 20). How the U.S. implements policy [Substack newsletter]. *Tech and Democracy*. https://carvao.substack.com/p/how-the-us-implements-policy

Carvao, P., Ancheva, S., Atir, Y., Zhou, B., & Jeloka, S. (2025, February). *Governance at a Crossroads: Artificial Intelligence and the Future of Innovation in America*. https://www.hks.harvard.edu/centers/mrcbg/publications/awp/awp251

Jerry McNerney. (2025, March 26). *McNerney Introduces Bill to Establish Safety Standards for Artificial Intelligence While Fostering Innovation*. Senator Jerry McNerney. https://sd05.senate.ca.gov/news/mcnerney-introduces-bill-establish-safety-standards-artificial-intelligence-while-fostering

Lima-Strong, C., & Hendrix, J. (2025, May 22). *US House Passes 10-Year Moratorium on State AI Laws | TechPolicy.Press*. Tech Policy Press. https://techpolicy.press/us-house-passes-10year-moratorium-on-state-ai-laws

National Science Foundation. (2025, February 6). *Request for Information on the Development of an Artificial Intelligence (AI) Action Plan*. Federal Register. https://www.federalregister.gov/documents/2025/02/06/2025-02305/request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan

Networking and Information Technology Research and Development. (2025, April 24). Comments Received in Response To: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan ("Plan"). *The Networking and Information Technology Research and Development (NITRD) Program*. https://www.nitrd.gov/coordination-areas/ai/90-fr-9088-responses/

The Datasphere Initiative. (2025, February 11). Sandboxes for AI. *The Datasphere Initiative*. https://www.thedatasphere.org/datasphere-publish/sandboxes-for-ai/

The White House. (2025a, January 23). *Removing Barriers to American Leadership in Artificial Intelligence*. The White House. https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/

The White House. (2025b, April 24). *American Public Submits Over 10,000 Comments on White House's AI Action Plan*. The White House. https://www.whitehouse.gov/articles/2025/04/american-public-submits-over-10000-comments-on-white-houses-ai-action-plan/