HARVARD Kennedy School
MOSSAVAR-RAHMANI CENTER
for Business and Government

Governance at a Crossroads: Artificial Intelligence and the Future of Innovation in America

Paulo Carvão
Harvard Kennedy School

Slavina Ancheva
Harvard Kennedy School

Yam Atir
Harvard Kennedy School

Shaurya Jeloka
Harvard University

Brian Zhou
Harvard University

February 2025

M-RCBG Associate Working Paper Series | No. 251

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

Governance at a Crossroads

Artificial Intelligence and the Future of Innovation in America

Paulo Carvão[1], **Slavina Ancheva**[2], **Yam Atir**[2], **Shaurya Jeloka**[3], **Brian Zhou**[4]

"What is past is prologue."
Shakespeare, W. (1610–1611). *The Tempest* (Act 2, Scene 1).

"If men were angels, no government would be necessary. If angels were to govern men, neither external nor internal controls on government would be necessary."
Madison, J. (1788). *Federalist No. 51*.

Abstract

The accelerated adoption of Artificial Intelligence marks a pivotal moment in technological progress. AI is reshaping industries, redefining labor markets, and prompting critical societal reflections on intelligence, reasoning, and the dissemination of information. While AI offers opportunities for economic growth, it also presents risks that must be managed to avoid adverse societal and geopolitical outcomes, making effective and transparent governance more urgent than ever.

This paper explores the potential of dynamic, collaborative public-private governance to foster safe innovation. Drawing from primary research, including interviews with tech industry leaders, U.S. Members of Congress, and staff, and an analysis of 150 AI-related bills introduced by the 118th U.S. Congress, this work identifies emerging areas of alignment between policymakers and industry stakeholders. It also highlights opportunities for a unified national approach, despite the challenges of a fragmented legislative environment.

The authors propose a dynamic governance approach that brings government and industry together while combining the foresight of ex-ante measures with the adaptability needed to respond to technological advancements. Coupled with existing ex-post mechanisms, the **Dynamic Governance Model** creates a comprehensive framework to promote competition, innovation, and accountability. It represents a policy-agnostic extra-regulatory framework, including a public-private partnership for standards setting and a market-based ecosystem for audit and compliance.

[1] First author - Mossavar-Rahmani Center for Business and Government, Harvard Kennedy School (paulo_carvao@harvard.edu)

[2] Co-second author - Harvard Kennedy School of Government (slavina_ancheva@hks.harvard.edu , yamatir@hks.harvard.edu)

[3] Technical contributor - Faculty of Arts and Sciences, Harvard University (shauryajeloka@college.harvard.edu)

[4] Primary Technical contributor - Faculty of Arts and Sciences, Harvard University (brianzhou@college.harvard.edu)

Ultimately, this governance approach can provide regulatory clarity and predictability, fostering an environment where businesses and innovation thrive while mitigating the risks inherent to AI's transformative power.

Keywords

artificial intelligence; AI; tech policy; innovation; industry; congress; Dynamic Governance Model

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 regulation. 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.

Slavina Ancheva is a Harvard Kennedy School Master in Public Policy (MPP) candidate and a Fulbright Bulgaria Scholar. She is a research assistant for Senior Fellow Paulo Carvão at the Mossavar-Rahmani Center for Business and Government. Before the Kennedy School, Slavina worked in the European Parliament as a Policy Adviser and Head of Office for one of the leading negotiators on the European Union's Artificial Intelligence Act, the world's first comprehensive law on AI.

Yam Atir is a Harvard Kennedy School Master in Public Administration (MPA) candidate. She is a research assistant to Senior Fellow Paulo Carvão at the Mossavar-Rahmani Center for Business and Government. Before attending the Kennedy School, Yam co-founded and led a tech policy think tank, focusing on macro-level research on technology ecosystems and providing strategic advice on global government-tech relations and emerging technology policies.

Shaurya Jeloka is an A.B. candidate in Computer Science and Applied Mathematics at Harvard University. He has developed intelligent algorithms for optimizing robotic systems, predictive models to mitigate M&A risks, and data-driven tools for performance analysis in academic and industry settings. Shaurya has also contributed to AI tools that enhance legislative research and policy analysis.

Brian Zhou is an A.B. candidate in Computer Science and Statistics at Harvard. He is a former researcher in autonomous systems at the Naval Research Lab and quantum machine learning at Brookhaven National Lab. Before Harvard, he founded a think tank focused on policy youth perspectives, cited by the Congressional Research Service and the UK Parliamentary Office of Science and Technology. He has testified at the UN Summit of the Future and helped establish the Ventures TECH program at the Harvard School of Engineering and Applied Sciences.

Table of Contents

List of Figures

Part I – The Context

A false choice

A storm was brewing unbeknownst to the captains of tugboats Montrose and T.J. Hooper off the coast of New Jersey, setting in motion a sequence of events that would reshape American tort law. In March 1928, when gale winds claimed their towed barges and cargo, few could have predicted the legal precedent emerging from these troubled waters. The subsequent litigation involved multiple parties: the cargo owners sued the barge operators, citing breaches of carriage contracts. Meanwhile, barge owners filed claims against the tugboat operators for unseaworthiness, specifically regarding their failure to implement radio receiver technology. This available safety measure could have provided advance storm warnings. Judge Learned Hand adjudicated the case (*The T.J. Hooper*, 60 F.2d 737 (2d Cir. 1932)), and his ruling established that the tugboat operators were negligent despite their adherence to contemporary industry standards. This verdict shaped a precedential framework, indicating that non-adoption of available safety technologies could constitute negligence, independent of industry-wide implementation rates. The ruling remained relevant in American tort law for nearly a century, informing legal discourse on technological innovation and safety standard requirements.

The debate around regulating technology is complex. Some argue that stricter regulation could hinder innovation by reducing incentives for investment and entrepreneurship. Others suggest that targeted regulation could promote innovation by incentivizing firms to focus on safety. These diverging perspectives are not new. We can learn from the past as we retrace the history of regulation and innovation. The Chief Executive Officer of a private, medium-sized, rapid-growth technology company shared his views about social media, another rapid-growth technology that AI often compares to. *"Social media is a weapon […] the most sophisticated propaganda weapon that was ever created […] The ability to tailor messaging, even before AI, is unprecedented. The pattern of usage of several times an hour is unprecedented. The combination of all of it is creating the most sophisticated way to brainwash people, and because of that, democracies collapse. […] We missed the train. We should have regulated social media."* The perspective of a senior tech investor, board member, and former operator highlights this complexity: *"People […] should have responsibility for the models that they create and govern those models across their life cycle, to understand what kind of risk profile the model represents […] being overly prescriptive on those things in this phase would be a mistake. […] Getting it right, in my mind, means not slowing down the pace of innovation because AI […] needs to progress."*

The current discourse on digital technology regulation and, specifically, artificial intelligence is missing the point. While the debate is often framed as a binary choice between whether to regulate or not, this is a false choice. Instead, we should focus on what regulatory approach will help steer innovation in a direction that our societal norms deem acceptable without stifling creativity, competition, and entrepreneurship.

Charting a path forward: Introducing a new governance model

How does one foster innovation while ensuring safety, fairness, and societal well-being? The answer lies not in a binary choice between regulation and laissez-faire innovation but in a balanced approach - a dynamic model capable of evolving with technology.

As the discussion of artificial intelligence regulation continues to evolve, policymakers face a dilemma: How can one foster innovation while ensuring safety, fairness, and societal well-being? The answer lies not in a binary choice between regulation and laissez-faire innovation but in a balanced approach - a dynamic model capable of evolving with technology. We propose a new model: an extra-regulatory framework[5] rooted in public-private collaboration, where regulatory clarity and accountability mechanisms are co-developed by government and industry.

The **Dynamic Governance Model** advocates for an incremental, iterative process. Rather than imposing rigid, one-size-fits-all rules, it emphasizes a responsive system of standards, audits, and compliance frameworks that adapt over time. The model encourages safe innovation while minimizing societal risks by fostering an environment where government oversight complements industry expertise. This approach acknowledges that today's faster pace of innovation requires a regulatory framework capable of evolving alongside technological advancements.

To guide readers through this multifaceted topic, we organized the manuscript into three parts:

1. **Historical Context and Technological Foundations**:

   The paper begins by exploring historical precedents in technology regulation and the unique characteristics of artificial intelligence. This part describes the AI triad - algorithms, data, and computing power - highlighting their role in AI's rapid advancement. It then introduces a new AI triad - energy, land, and labor. The first part concludes with a discussion about the attributes of AI policy, examples from around the world, and the specific aspects of how the United States implements policy.

2. **Current Legislative Landscape and Industry Perspectives**:

   Following the initial context setting, the paper presents a first-of-a-kind analysis of the AI proposed legislation during the 118th United States Congress Session. The analysis is enriched by insights from qualitative interviews conducted by the authors with Members of Congress and their staff and industry leaders. The interviews highlight emerging areas

[5] A word on the term "extra-regulatory framework": This work does not advocate for any specific legislation or statute to regulate the technology. The focus, instead, is on a method that can be applied to different policy objectives, from smaller, niche-based policies to larger AI national policy initiatives.

of consensus between policymakers and the private sector, as well as key gaps where greater alignment is necessary. These findings form the foundation of the proposed governance model.

3. **Proposed Framework for Dynamic Governance**:

   The final part of the paper outlines the Dynamic Governance Model's three core components:

   a. Public-private partnerships for the creation of evaluation standards.
   b. A market-based solution for audit and compliance.
   c. A system of accountability and liabilities set by legislatures, existing executive agencies, and the courts.

Our objective is to provide an actionable path toward balanced AI governance, first establishing the historical and current context and then introducing a forward-looking solution. While the reader should approach these three parts sequentially, one can jump directly to the segment of their particular interest and only refer to others as needed.

The renowned American legal historian and scholar, Lawrence M. Friedman, said: "[T]echnology and science, as they move and change, impact the general culture. The general culture then, in turn, impacts law, the legal system, and the very structure of government" (Friedman, 2023). Let's start at the beginning of the quest for Artificial Intelligence.

A brief history of artificial intelligence

It was as if the dreams devised during the summer of 1956 had suddenly become a reality. In November of 2022, sixty-six years after a group of mathematicians and scientists got together for the Dartmouth Summer Research Project on Artificial Intelligence (McCarthy et al., 2006), OpenAI's release of ChatGPT captured the zeitgeist by placing a mirror in front of society. The imitation game proposed by Alan Turing was on, and his question, "Can machines think?" was being asked by the masses (Turing, 1950).

Figure 1 OpenAI's CEO Tweet Announcing ChatGPT's Release to the World

The 1956 summer artificial intelligence research, considered the Constitutional Convention of AI ("Encyclopedia of Artificial Intelligence: The Past, Present, and Future of AI," 2021), assumed that one could precisely describe all aspects of human learning and other intelligence features so that a machine could simulate it. Based on that premise, the researchers who met at Dartmouth sought to discover how machines use language and form abstractions to solve problems typically addressed by humans. Thus, the race started with the expectation that they would make substantial progress over a few weeks of dedicated effort. The progress, however, was slow, with scientific discoveries followed by a series of long hiatuses – the AI winters - in which academic research and industry focus drifted away from the topic. Fast forward to 2022, ChatGPT reached 1 million users in just 5 days after launch, and as of October 2024, it has approximately 250 million active users (*ChatGPT Revenue and Usage Statistics (2024)*, n.d.) To understand this dynamic, we must look closely into how AI works and the drivers of its progress.

AI stack – algorithms

One can think about modern artificial intelligence as intelligent mathematical algorithms being trained on massive amounts of data, using, in this process, unprecedented computing power. This combination of algorithms, data, and computing power is commonly known as the AI triad (Buchanan, 2020). The history of the current AI algorithms started in the 1940s when a neurophysiologist and a logician modeled neural activity via logical operations, establishing the theoretical basis for what we today call neural networks (Warren S. McCulloch & Walter Pitts, 1943). More than a decade later, using the then-current state-of-the-art IBM computers at the Cornell Aeronautical Laboratories, Frank Rosenblatt, a psychologist, worked on the perceptron, an early version of binary classifier algorithms inaugurating the field of deep learning (a form of machine learning which imitates how the human brain functions requiring training with large amounts of data) (Rosenblatt, 1958).

During the next couple of decades, scientists struggled with how to scale and perfect the training of neural networks. This lasted until the late 1980s when scholars introduced the concept of backpropagation, enabling the efficient training of multi-layered neural networks (Rumelhart et al., 1986). This breakthrough allowed the development of the first convolutional neural network architecture, which was initially applied to handwriting recognition and designed to mimic the visual processing of the human brain (LeCun et al., 1989). During the 1990s, Recurrent Neural Networks drew attention given how they process sequential data and maintain the memory of past input. These techniques were applied in areas like speech recognition and natural language processing (Hochreiter & Schmidhuber, 1997). In 2012, researchers at the University of Toronto won the ImageNet Large Scale Visual Recognition Challenge using deep convolutional neural networks and showed the promise of deep learning for computer vision (Krizhevsky et al., 2012). The stage was set for the introduction of the Transformer architecture by a team of Google and University of Toronto scholars in 2017. This streamlined design eliminated the need for recurrence and convolutions, leading to superior training quality while maximizing parallel computing and significantly reducing training time (Vaswani et al., 2023). Shortly after that, in 2018, scientists at OpenAI introduced the Generative Pre-trained Transformer (GPT), a large language model (LLM)

that uses neural networks to understand and generate text (Radford, 2018). The Generative Pre-trained Transformer serves as the core algorithmic architecture for ChatGPT.

Recently, scientists and engineers have focused on energy consumption optimization, improving problem-solving capabilities, and developing models that can operate across multiple domains beyond language with increased agentic features. Techniques like Mixture of Experts (MoE) (Bergmann, 2024), Chain of Thought (OpenAI, 2024), and multimodal agentic offerings (Pichai et al., 2024) are hitting the market. Researchers from the University of California Berkeley recently have open-sourced Sky-T1, a reasoning AI model fine-tuned for under $450, enabling affordable advanced AI development. However, “training” here means fine-tuning, not pre-training, which remains resource intensive (Wiggers, 2025).

This summary is not an exhaustive list of all the achievements and pioneers in this field; instead, it provides a bird’s-eye view of some of the key milestones in the development and evolution of AI algorithms.

Two distinct timelines emerge when reviewing this brief history of artificial intelligence algorithms: One from 1943 to 2018, seventy-five years of relentless work creating, tweaking, and optimizing the mathematics and computational structures. Another was in the seven years from 2018 to 2025, particularly since November 2022, when ChatGPT exposed the transformer architecture and GPTs to the public via an easy-to-use chat interface. What makes them different? Data and computing power, the other two elements of the AI triad, will provide the necessary clues.

AI stack – data

Data has been an integral part of computer science, and the need to manage large data sets goes back to the early data centers used by governments and banks in the 1960s and 1970s. Relational databases, invented at IBM research labs in 1970 (Codd, 1970), were used for most traditional transactional processing. As we entered the new century, the rate and pace of data generation increased. Internet usage and e-commerce became pervasive as more intelligent devices connected to the web. We have transitioned from an era where only computers generated data to one defined by the Internet of Things, with data now produced by home appliances, cars, phones, and a wide range of other devices. Social media connects billions of people worldwide, constantly using their mobile devices and generating even more data. Data became ubiquitous, and the costs for generating, transmitting, and storing it were reduced drastically (Mathieu, 2024).

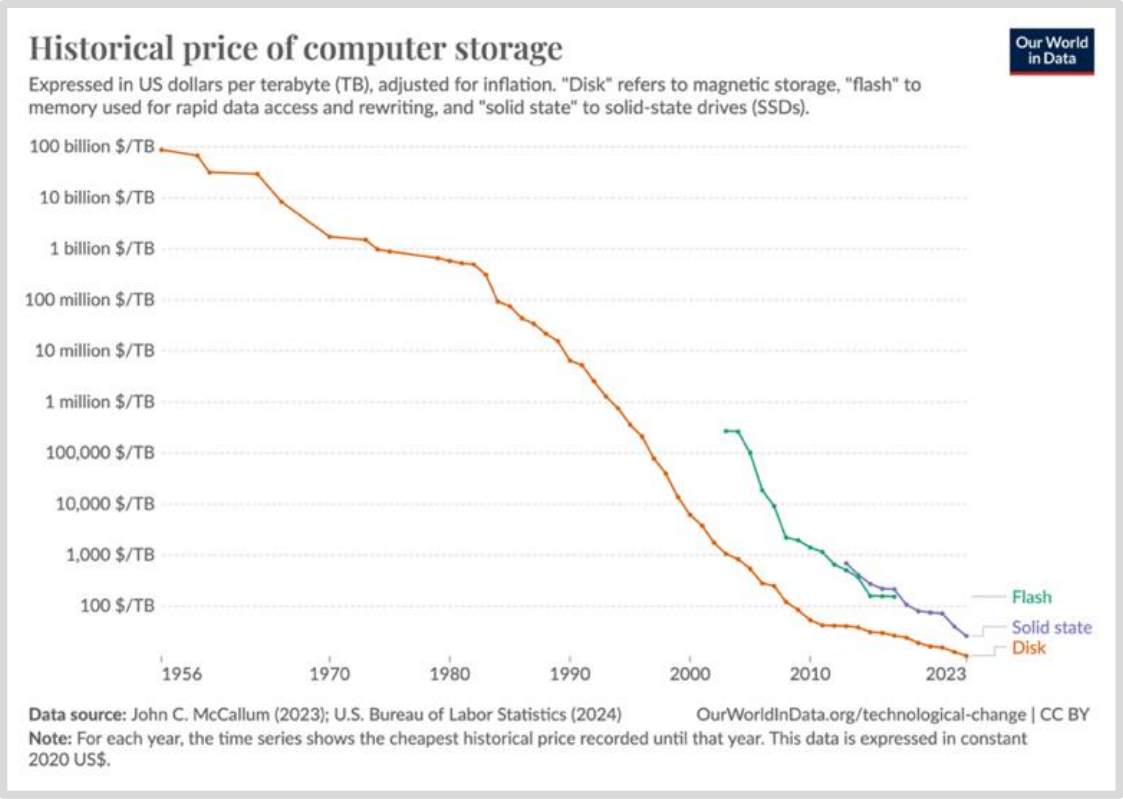

Figure 2 Decline of Computer Storage Costs - Our World in Data is an effort between the University of Oxford and the non-profit Global Change Data Lab.

AI stack – computing power

Having discussed two of the three components of the AI triad - algorithms and data – let us look at computing power. Here, two phenomena played a decisive role: Moore's Law, which observes that the number of transistors on an integrated circuit doubles approximately every two years while maintaining a relatively constant cost (Moore, 2006), and the advent of cloud computing. To put this into perspective, Intel's first microprocessor, the Intel 4004 from 1971, had 2,300 transistors (Ars Technica, 2011). An NVIDIA Blackwell Series GPU announced in 2024 with planned availability in 2025 will pack 208 billion transistors (*NVIDIA Blackwell Architecture*, n.d.). Cloud computing is delivering computing services over the Internet, doing away with the need for proprietary data centers. The concept of sharing computing capacity goes back to the 1960s and 1970s when time-sharing and virtualization of traditional IBM mainframe computers were used. In its modern form, cloud computing has taken off since 2002 with the launch of Amazon Web Services (AWS), followed by Google Cloud Platform in 2008, Microsoft Azure in 2010, IBM SmartCloud in 2011, and Oracle Cloud in 2012. Cloud deployments reduce the initial investment needed to build computing capacity and allow for flexible consumption "by the drink."

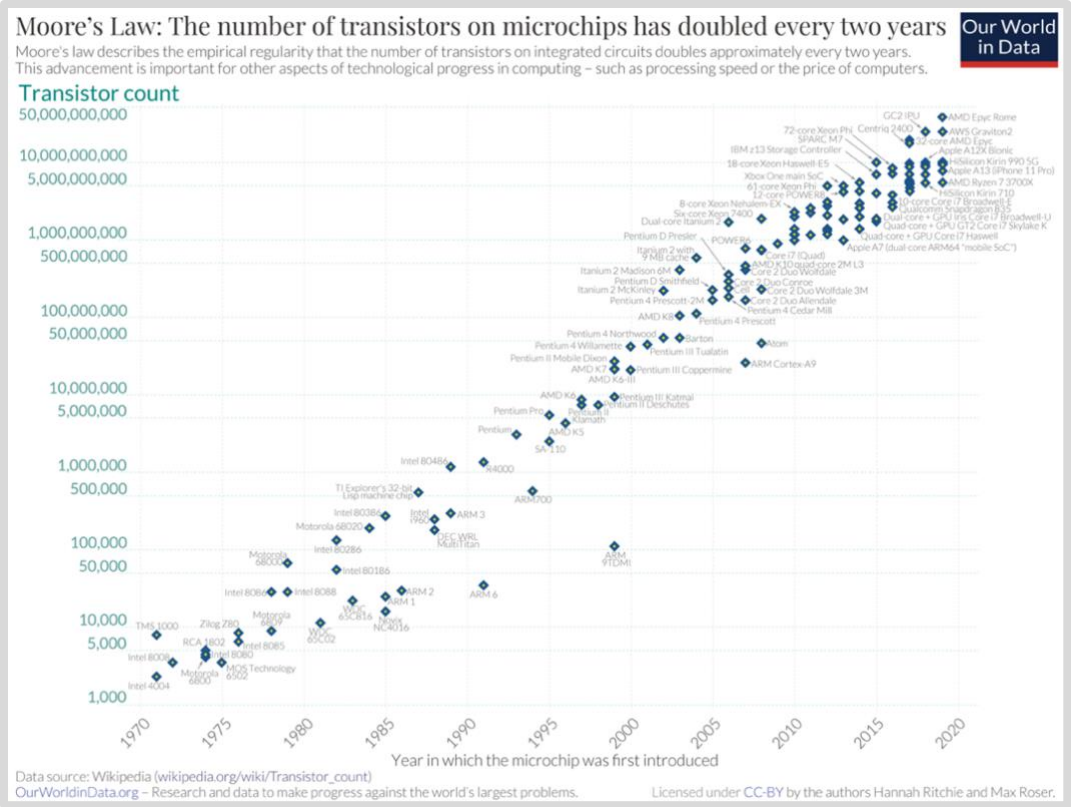

Figure 3 Increase of Semiconductor Density (number of transistors on chips) - Our World in Data is an effort between the University of Oxford and the non-profit Global Change Data Lab.

This completes the triad: modern algorithms, with access to large training data sets, and leveraging lower-cost flexible computing power. The confluence of these three factors explains the acceleration since 2018.

Large Language Models have become the public face of AI today, accessible to users through chat interfaces or APIs. Examples are ChatGPT from OpenAI, the Doubao family of models from ByteDance, Gemini from Google, Claude from Anthropic, DeepSeek-V3 and R1 from DeepSeek, Grok from xAI, LLaMA from Meta, Granite from IBM, Qwen from Alibaba, Mistral Large from Mistral AI, ERNIE from Baidu, Hunyuan from Tencent and more. Since their initial introduction, material progress has been made, with most advances driven by scaling — the ability to process increasingly large training datasets (Kaplan et al., 2020). The size of a Large Language Model (LLM) is typically measured by the number of parameters it contains, which has been growing exponentially. For instance, ChatGPT-3 has 175 billion parameters, while ChatGPT-4 contains 1.7 trillion parameters. However, this scaling comes at a significant cost, as the energy consumption required to train and operate LLMs raises serious concerns about their environmental impact (Carvao, 2024b). This rapid growth, combined with unresolved technical challenges such as transparency, explainability, and the mitigation of hallucinations, has prompted some scholars to question whether LLMs are approaching their limits (Marcus, 2024). These conversations echo in the halls of Congress and state legislatures, where decision-makers

oscillate between *"Wow, AI is coming, and we're on the cusp of superhuman intelligence"* and *"We're hitting a plateau."* Ultimately, *"both things can be true… [and] there's no way we're going to keep a pulse on everything all at once,"* concludes one Staffer.

The second AI triad – energy, land, and labor

AI and its data center infrastructure are hungry for energy.

According to Goldman Sachs, data center power demand will surge to 8% of the U.S. total by 2030 from 3% in 2023, growing at a 15% CAGR (compound annual growth rate). For context, the total energy demand in the country has been flat in the decade up to 2022 and is now projected to accelerate at 2.4% CAGR, of which data centers represent 37.5% of the total growth. This pattern reflects an AI-driven growth in data center demand at the same time there is industrial reshoring and growth in manufacturing activity. About $50B in capital investment in U.S. power generation is required only to keep up with the data center growth until 2030 (Davenport et al., 2024). In another projection, the U.S. Department of Energy estimates that data centers will consume up to 12% of total U.S. electricity by 2028 (U.S Department of Energy, 2024). This pressure on the country's power system has caught the attention of Congress as reflected by this congressional staffer's comments: *"A report […] out yesterday from Lawrence Berkeley […] it's going to take up 12% of our electricity demand […] it's going to have a higher impact on the new electricity demands on the grid, more than EVs and onshoring and electrification of processes. […] And as a result, something needs to be done, whether that be engaging the AI companies […] to build out the systems that we need to power what they're doing, or that is using these AI systems in a way to solve these problems. […] real people are suffering from rate hikes."*

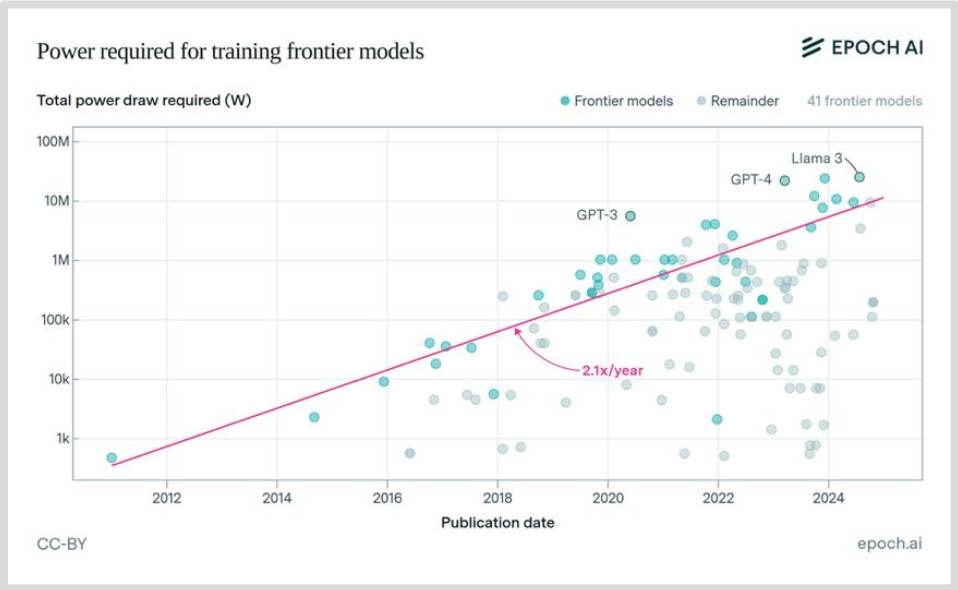

Figure 4 The power required to train frontier models is doubling every year (Epoch AI, 2024)

Big Tech is accelerating investments in AI data center buildout. Microsoft announced plans for a 2025 investment of $80B in data centers to train AI models and host the cloud infrastructure to deploy AI applications. While this represents a global investment, more than half will be made in the U.S. (Smith, 2025). OpenAI, together with Oracle, SoftBank, and the MGX Emirati investment firm, announced a $500B investment over four years to build AI infrastructure (OpenAI, 2025). The AI-driven data center boom is reshaping several U.S. regions. Northern Virginia's "Data Center Alley" leads, consuming a quarter of the state's power load. Texas follows with rapid growth, while California remains crucial despite constraints. Washington state has emerged as a significant player, with data centers consuming 5.69% of its electricity in 2023. The state benefits from abundant renewable energy, with 54% of Central Washington's data centers powered by hydroelectricity. However, this growth strains the power grid and challenges Washington's green energy goals. Emerging secondary markets are being developed across the country, bringing a combination of economic benefits and infrastructure and environmental challenges as states balance growth with sustainability concerns (Cheung, 2024; Tremayne-Pendelly, 2024).

The surge in demand is driving large enterprise consumers to reevaluate their energy-sourcing strategy. It has led to a resurgence of interest in nuclear power for reliable and emissions-free electricity. One technology, small modular reactors (SMR), has caught the attention of Big Tech. SMRs' compact footprint allows them to be sited in locations unsuitable for larger nuclear plants. Their prefabricated units can be manufactured off-site, shipped, and installed efficiently, reducing costs and construction delays compared to custom-built large reactors. These smaller reactors can be deployed incrementally, aligning with growing energy demands while saving time and expense (Liou, 2023). Amazon Web Services incorporated SMRs into their plans for net-zero carbon operations by 2040 and is signing agreements with utilities in Washington and Virginia to use this technology (Amazon, 2024). In addition to exploring the usage of SMRs, Microsoft plans to reactivate the Three Mile Island nuclear power plant in Pennsylvania, the site of the worst nuclear reactor accident in U.S. history (Eccles, 2024; Plumer, 2024).

Google is collaborating with the startup Kairos Power to develop and deploy SMRs. Their goal is to provide up to 500 megawatts of continuous carbon-free power by 2035 to support their AI energy demands while helping improve grid reliability and global decarbonization efforts (Terrell, 2024). A congressional staffer questions whether we can get there fast enough: *"I'm very pro-nuclear. I think that we should be getting more nuclear on the grid as fast as possible. I'm also very aware that SMRs are unproven […] and that the larger reactors, like Three Mile Island, if we want to get it back up, will take years, and the ones we want to build from scratch will take decades. […] nuclear is going to be a really great long-term play, but something needs to be done in the short term too […] the issue is no one, I've heard, no one at any level in any country, has a short-term play."*

In parallel, in an intense lobbying effort, OpenAI has proposed to the Biden and Trump administrations the development of massive data centers, each consuming as much power as an entire large city, to advance AI models and maintain U.S. competitiveness with China. The company argues that these facilities could generate tens of thousands of jobs and boost GDP if supported by the right governmental policies. The proposed 5-gigawatt data centers each

represent the equivalent of five nuclear reactors and could power 3 million homes (Ghaffary, 2024). *"There's one-gigawatt, there's two-gigawatt, there's three-gigawatt data centers that are being planned right now and are being talked about. That's more energy than most cities need, like small cities, it's absurd. So, you know, it's a thing that we really need to be thinking about because every time one of those goes on the grid, it makes the grid less reliable"* highlights our interviewee. Oklo, a U.S. nuclear power plant developer, has signed a non-binding agreement with data center operator Switch to deploy 12 gigawatts of projects based on their technology by 2044 (World Nuclear News, 2024). Sam Altman, CEO of OpenAI and Chairman of Oklo's Board of Directors since 2015 (Oklo Inc., n.d.), actively lobbied the White House for government support in the nuclear-powered data center buildout. This episode underscores the promise of public-private partnerships and policies to drive technology adoption while testing the boundaries of conflicts of interest when a single company dominates multiple aspects of the discussion.

The new facilities require repurposing historically agricultural land to develop large-scale data centers. Data centers rely heavily on water for cooling, often drawing hundreds of thousands of gallons per day and putting considerable pressure on natural resources (Carvao, 2024b). Orchards give way to high-tech hubs, disrupting local economies and cultures. The influx of wealth builds schools and infrastructure but raises living costs, pricing out many long-term residents. These dynamics impact the real estate market, with developers now seeking much larger land parcels, often 500 to 1000+ acres, compared to the 15–20-acre plots sought a decade ago. The scarcity of suitable sites with adequate infrastructure pushed development into secondary markets and rural areas. Developers are facing resistance from local communities concerned about the impact of large data centers on their areas, and some counties are considering bills to curb data center development. Land prices in key markets have tripled while asking rates for data center space have risen by 20% year-over-year in major North American markets (CBRE, 2024; Houlihan Lockey, 2024).

Data center construction across the United States created a surge in demand for skilled electrical trade workers to build facilities and infrastructure. The current lead time to power new data centers in large markets such as Northern Virginia can be more than three years. And, in some cases, lead times for electrical equipment are two years or more. This escalating demand underscores the need for substantial investments in workforce development to ensure the availability of qualified labor essential for the timely and efficient expansion of data center infrastructure (Green et al., 2024). Labor constitutes the backbone of AI's ascent. Migrant electricians work long hours for lucrative pay while unions scramble to train new workers. Yet, once operational, these facilities offer minimal jobs, underscoring the short-term nature of such booms.

This second AI triad of energy, land, and labor reshapes economies and landscapes. Central Washington illustrates this transformation, where abundant hydropower fuels sprawling data centers vital for AI. Power drives the infrastructure as small towns become magnets for data-intensive operations. However, this reliance stretches the grid, forcing companies to seek nuclear solutions (Weise & Tamayo, 2024).

AI data centers' increasing electrical power demand has a significant environmental impact. The training of GPT-3 by OpenAI used 1,287 megawatt-hours of electricity and produced 502 tons of CO2 emissions (Czarnecka, 2024), the same power 120 average American homes consume in one year (Aibin & Simic, 2024). Using LLMs for search is becoming the norm. Directly, with an app like Perplexity, using ChatGPT's search feature, or relying on Google's AI-generated search summaries. This new way of searching, at inference time, uses as much as ten times the energy a Google search does (Goldman Sachs, 2024).

The new AI triad - **energy, land, and labor** - illuminates a new series of questions to be addressed by policymakers.

The new triad illuminates a new series of questions to be addressed by policymakers. Another congressional staffer working in the Senate told us: *"We want to be strategically competitive against China, and so we want to have all the AI resources we possibly can. That's a question of IP and foundation models. It's a question of chip resources. It's a question of energy availability to power AI models, and I think these things […] give us the upper hand."*

The new AI triad - energy, land, and labor - underscores substantial opportunities for public-private partnerships while highlighting the need for policy development or reform to support industry growth and address challenges. Setting aside the merits or politics of each area, issues such as land permitting reform and workforce adaptation and training emerge as priorities. The congressional staffer continues: *"Both parties can come together and say, well, we cannot build AI data centers on solar and wind […] our current infrastructure is totally insufficient to support the power needs of AI centers across the country. But there exists new technology […] if you can release some of the regulations so that they can be built in five or six years and not 14 or 15 years […]. Energy infrastructure is not a short-term play, and it has to survive across Congresses and across administrations. […] China plays the long game. China understands that this is not a short-term play. This is a long-term play, but the U.S. is always trying to get the best bang for his buck in four years."*

There are stark contrasts between the rapid evolution of industry demands and the slower pace of legislative processes. Reconciling these two dynamics will be a recurring theme in our discussion.

A unique moment?

Here we are, witnessing a technology undergoing rapid adoption while unresolved questions remain (Lynch, 2024; Strickland, 2024).

Is this AI moment truly unique? In many ways, the recent rise of artificial intelligence adoption is exceptional, especially considering that it follows the slow burn of the first seventy-five years of science in the field described in a previous section. The recent advancements in AI can be attributed to the cumulative effects of progress in microelectronics and the more effective

deployment of neural networks - concepts that have existed for decades. While AI, like many other elements of the technology and venture-backed industries, is often overhyped, its impact on personal and business productivity is starting to become visible. Previous automation waves have primarily affected manual and clerical work. AI, especially when combined with other forms of physical automation, such as robotics, can also have a material impact in this area. In addition, AI can transform white-collar and academic work, especially for repetitive, well-structured tasks, and its implications for knowledge workers are different than what previous waves of technology-driven automation brought (Dell'Acqua et al., 2023).

A recent working paper from the National Bureau of Economic Research explores the role of Generative AI as a new general-purpose technology. The article investigates its rapid adoption, using a nationally representative survey in the U.S., and finds that 39% of adults aged 18–64 use Generative AI, with 28% using it for work tasks like writing, data analysis, and coding. Adoption is higher among younger, educated, and higher-income individuals, with significant use in management and technical occupations. Generative AI has been adopted faster than prior technologies like PCs and the Internet, driven by its consumer-oriented nature. While adoption varies by demographics and occupation, early patterns suggest productivity gains, though its broader economic impact remains debated (Bick et al., 2024).

On the other hand, this is not the first general-purpose technology to have a deep and broad impact on economies and societies. Policies and structures have been established during other waves of technological change to stimulate and safeguard adoption while mitigating and addressing risks and harms. As one of our congressional interviews highlighted, the period in which *"everyone collectively freaked out about AI last year is starting to come down from the hype cycle."* Some industry leaders echo this sentiment: *"We didn't have enough of a regulatory regime for Blockchain and Web 3, and we had a massive failure in FTX […] governments are afraid of missing that again, and the harm to investors and others that came from that is very fresh right now. There's also a feeling of the uniqueness of these technologies and systems […], but I'm, I'm not so convinced. I don't believe you need a quantum regulatory regime that would be materially different from an AI regime."*

Nevertheless, unresolved challenges, such as the development of autonomous agents, remain on the horizon. It is imperative to look at both sides of this argument.

Throughout the history of scientific and technological progress, the explosive growth and accelerated adoption of AI stand out as a distinct moment. AI is reshaping multiple facets of our lives — from the organization of production and labor to the balance between automation and job creation, the control and dissemination of information, and even the very concept of intelligence and reasoning (Acemoglu & Restrepo, 2020; Harari, 2024). These transformations carry profound socio-economic and geopolitical implications that often surpass the impact of previous technological shifts. A Senate staffer with deep technological experience put it in vivid terms: *"I think that artificial intelligence will prove to be the most influential technology discovered by man since the discovery of the wheel, the fire, and the Internet in terms of the impact it will have on*

the productivity and the survivability of mankind. […] Anyone who thinks it won't be is not aware or doesn't understand the ultimate potential."

The recent rapid pace of AI development, its ability to operate autonomously, and its widespread impact across multiple sectors make it fundamentally different from other technologies. These increasingly complex systems can make decisions without human intervention (Chan et al., 2023), raising questions about accountability, as traditional liability frameworks are designed around human actions. For example, if an autonomous AI system makes a harmful decision, it is unclear who is responsible: the developer, the deployer, or the AI system itself. The emerging agentic nature of algorithmic systems is one of the main ways they differentiate from traditional technologies. There is a regular interplay between AI technology and societal values like fairness, privacy, and transparency. For instance, biased algorithms, or algorithms trained on biased data, can perpetuate discrimination, and AI in surveillance can infringe on privacy. As previously discussed, AI systems rely heavily on data, which raises privacy and security concerns as training datasets can include personal or sensitive information. AI systems often operate as "black boxes" where even their creators cannot fully explain how they arrive at specific outcomes. In many ways, **AI is unique**.

But is it? Looking back at history, one can identify major transitions in economic development from the beginning of organized economic activity to today. While there are multiple ways to categorize these phases, let's look at the defining emerging technology from each era. The Neolithic Revolution (c. 10,000 BCE) marked the transition from hunter-gatherer societies to agricultural ones. The systematic cultivation of plants and the domestication of animals acted as the defining technologies. The Urban Revolution (c. 3000 BCE) saw the rise of the first cities and complex societies enabled by the development of written language and record-keeping systems. During Classical Antiquity (c. 800 BCE - 500 CE), the development of more sophisticated economic systems was enabled by extensive systems of roads, particularly the Roman roads. Later, a transition from feudal to commercial, market-oriented economies (c. 1000 - 1500 CE) happened in parallel to the advent of the printing press, which revolutionized the spread of information and knowledge. The First Industrial Revolution (c. 1760 - 1840) was a transformative period that saw the shift from agrarian economies to industrial ones and the introduction of the steam engine powering factories, trains, and ships, mechanizing production, transforming transportation, increasing urbanization, and ushering the rise of factory systems and wage labor. The Second Industrial Revolution (c. 1870 - 1914), built upon the first, introduced mass production and new industries (like chemicals, steel, and petroleum), gave birth to large corporations and scientific management, and was powered by electricity. The post-industrial phase sometimes referred to as the information or knowledge economy (c. 1950 – present), occurred as manufacturing gave way to service-based economies, with the entrance of information technology and digital economies, the globalization of production and finance, and the emergence of knowledge-based industries. Semiconductors, computers, and the Internet revolutionized productivity, communication, and commerce.

Each of these transitions fundamentally altered the structure of economic activity, the organization of labor, and the distribution of wealth and resources. They represent key inflection points. The

underlying technologies in each era were transformative. These General Purpose Technologies, such as the steam engine, the electric motor, and semiconductors, became pervasive across a large cross-section of the economy, facilitating widespread productivity gains. General purpose technologies drive "innovational complementarities," and the productivity of a downstream sector increases with the innovation introduced with the enabling technology. The development of downstream applications increases the return to advances in the enabling general-purpose technology. Advances in this technology lead to opportunities for new applications in a positive feedback loop of accelerated technical progress and economic growth (Bresnahan & Trajtenberg, 1995). The current adoption of AI mirrors past instances of technological change, where new possibilities can disrupt social, environmental, and cultural values, highlighting the broader challenge of law struggling to adapt to the rapid pace of innovation and its societal impacts (Bennett Moses, 2007). From this historical perspective, **AI is not unique** and may only be another general-purpose technology unleashing growth through productivity. The pace of change, however, is accelerating.

We are amid an emerging new phase in a post-knowledge economy. There is an increased emphasis on uniquely human abilities such as creativity, emotional intelligence, ethical reasoning, and complex problem-solving. The youth seek purpose and meaning, and economic activity focuses on individual and societal well-being rather than just optimizing productivity. Adaptability and resilience are at a premium as citizens adapt and respond to future challenges. Traditional knowledge-based tasks are beginning to be taken over by artificial intelligence, freeing humans to focus on higher-level creative and meaning-oriented work. AI may become the defining technology of this new era, but it is too early to tell.

There is a growing body of scholars questioning economic growth (Gross Domestic Product or GDP) as the primary measure for success at the expense of human well-being and environmental health. They argue instead for balancing human needs with planetary limits by adopting economic development that is regenerative by design (Raworth, 2017). This perspective can be contrasted with the techno-optimistic view embraced by some leading investors who are funding venture capital-fueled innovation (Andreessen, 2023; Dean, 2023). As in other dual-use technologies, there are plenty of opportunities for malicious use of AI with the capacity to weaponize it in ways that previous technologies like nuclear, chemical, and biological have been (Herdman, 1993). One of our industry interviews captured this dichotomy: *"There's still a huge, huge opportunity ahead, and there has been tremendous progress in recent years, both with the technology as well as the applications and the use cases. Broader adoption has allowed us to accelerate our understanding of what AI could be used for and how to use it. And as technologies […] have gotten into the hands of so many more people, now we're seeing feedback loops and improvement at a rapid rate. We're seeing exciting and promising use cases as well as concerning use cases. […] I'm a technologist, I'm an early adopter, I'm somebody who believes in the promise of technology. At the same time, I don't think technology is inherently good or bad, but we need to have appropriate guardrails and […] a tech policy discussion."*

While researchers continue to push the boundaries of the AI triad, early signs of saturation in data availability and the growing energy demands of large-scale computation pose significant

challenges. Scientists use synthetic data to augment training datasets and reduce reliance on real-world data to address these limitations. In parallel, model distillation techniques are used to create smaller, more efficient AI models requiring less computational power while maintaining high performance. For instance, a large LLM trained on real and synthetic data can be distilled into a more compact version that runs on edge devices, making AI more accessible and sustainable. As recently as 2023, authors wrote about model deterioration when training relies on increasing amounts of synthetic data (Shumailov et al., 2023). As one begins to approach the limits of legally accessible data on the Internet, the use of machine-generated data becomes a critical path forward. A team of researchers at Microsoft recently published a new model with performance comparable to larger ones while utilizing high-quality synthetic datasets in combination with high-quality organic data and employing post-training innovations (Abdin et al., 2024). Progress using synthetic data is one of many examples of the interplay of innovation and progress.

Generative AI exposed new frontiers in accessing, manipulating, and generating information. We saw the first round of the impact on knowledge workers and professions. As we move beyond the first quarter of the century, we are entering an era of agentic AI, delivering on the promise of embedding intelligence into business processes. Generative AI agents are evolving from information providers to autonomous actors capable of executing complex workflows and collaborating seamlessly with humans, promising enhanced productivity and innovation (Yee et al., 2024). AI agents have autonomous goal-driven behavior, contextual decision-making, and iterative learning. Unlike static models, agents dynamically interact with environments, leveraging reinforcement learning and LLMs to perform complex, multi-step tasks without human intervention. Engineering excellence and technical innovation have once again unleashed the agentic AI era we are embarking on (Wiesinger et al., n.d.). Reduced human involvement and increased autonomous decision-making heighten the need for ethical development and awareness of engineers' sociotechnical impact.

Current challenges like model footprint, the associated computing power and energy consumption, model transparency, hallucinations, alignment, and more can and must be remedied by creating new algorithms and techniques to address the shortcomings. Over time, the industry can develop and deploy inherently safe models with minimized externalities, provided the right incentive policies are in place. This approach can sustain the current pace of technological advancement and adoption.

Charting a path forward requires thoughtful reflection on AI's unique - or perhaps not so unique - nature. It demands a careful balance between the utopian visions of techno-optimists and the dystopian concerns of those who question whether our growth-driven economic model places an unsustainable burden on the planet. While AI possesses distinctive characteristics and is advancing at a pace that challenges our capacity to adapt, technological change is not a new phenomenon. Lessons from the past can offer valuable insights.

While AI possesses distinctive characteristics and is advancing at a pace that challenges our capacity to adapt, technological change is not a new phenomenon. Lessons from the past can offer valuable insights.

From a policy perspective, the evolution of AI is a blend of continuity and novelty. AI is often compared to social media by several elected representatives. *"I think one of the mistakes we made with the Internet and social media is that we were too late to the game, and we lacked putting in the proper guardrails,"* says one elected representative. Yet, despite this consensus that Congress may have missed an opportunity to act on social media, the parallel has its limits. *"The difference is social media doesn't have a super positive use case – connecting people is meaningful – but AI innovation will change defense and competition with China,"* says one congressional staffer

Most of our interviews with Congress distinguished between AI uses already addressed by sectoral legislation and the emergence of Generative AI, which may warrant a distinct regulatory approach. For instance, using AI to commit fraud or discrimination is not unique, as these behaviors can be governed through existing sectoral laws. A congressional staffer summed it up best: "*If someone uses a pencil or an AI technology to carry out discrimination, they will have to deal with that discriminatory act*." At the same time, unlike most traditional policy issues, Generative AI and general-purpose AI more broadly cut across multiple domains. As another congressional interview highlighted: *"One [issue] is how do you regulate AI products with well-defined use cases. The other is how to regulate general-purpose AI that can be used for lots of things […] Currently, we don't know what a meaningful framework [for general-purpose AI] looks like, but we definitely need one… it's difficult."*

This dilemma has led to a proliferation of bills introduced by Members of Congress from various committees, parties, and interest areas. This effort began as early as the 116th Congress but picked up steam after the introduction of ChatGPT, as evidenced by the volume of AI-related legislation introduced in the 118th Congress compared to prior sessions. A notable trend from the early days of AI policymaking was the emphasis on education for Members of Congress and their staff. Congressional caucuses, AI task forces, and the formation of congressional Bipartisan AI Caucuses played a key role in equipping lawmakers with the knowledge required to craft informed AI policy. Similarly, committee staff members see education as a core component of their role, ensuring policymakers are equipped to address the rapidly evolving AI landscape.

AI policy considerations

"The marvels of technological advance are not always risk-free", a renowned legal scholar reminds us (Mandel, 2007). Today's literature extensively covers the benefits and risks of artificial intelligence. AI offers benefits such as increased productivity, personalized services, and advances in healthcare, education, and decision-making. Automation reduces human error, streamlines operations, and enables continuous production. However, AI also carries risks, including labor displacement, data privacy violations, algorithmic bias, and the concentration of power among a few corporations. It can be exploited for mass surveillance, disinformation, and manipulation of public opinion. Autonomous weapons and AI-driven cyberattacks add security threats. System failures and unpredictable AI behavior create further vulnerabilities. This duality underscores the need for careful policy considerations (Barroso, 2024; Carvao, 2024a).

Efforts to establish technology policy balance a diverse set of attributes: elements of market power and concentration (typically associated with antitrust), safety and responsible use of technology, aspects of industrial policy, national security, intellectual property, environmental issues, and concerns with the public interest. This list is representative, albeit non-exhaustive. The overarching policy goals and normative values in each society/jurisdiction will determine the relative importance of each element.

No specific configuration is better than another; what is important is to achieve the policy goals and ensure that the mechanisms used to that extent are enforceable. In addition, a good policy should enable and protect innovation and benefit businesses, the economy, and society. This balance is not straightforward; different approaches can be taken for each policy attribute. A congressional staffer highlighted: *"Some regulations are good; most regulations are probably good. People have a lot of opinions about government bureaucracy and inefficiency, which are valid and […] that's a hard thing to sell to the public or anyone in a decision-making capacity."* On the industry side, a mid-sized company's chief executive provided a candid and direct feedback: *"unless the federal government figures out what they're doing sooner rather than later, it's going to be a mess […] if you start having all these hodgepodge of state regulations, it's just another mess to manage making sure you fit all these regulations."*

Policy goals will determine the relative importance of each strategic attribute. No specific configuration is better than another; what is important is to achieve the stated goals via enforceable mechanisms.

Elected governments should set the goals.

Figure 5 Policy Attributes

Until 2024, no federal laws in the United States govern artificial intelligence, except for the non-regulatory National Artificial Intelligence Initiative Act of 2020, included in the 2021 National Defense Authorization Act during the first Trump administration (*Text - H.R.6216 - 116th Congress (2019-2020)*, 2020; *Text - H.R.6395 - 116th Congress (2019-2020)*, 2021). The law established a federal initiative to accelerate AI research and development, created the National Artificial Intelligence Initiative Office under the White House Office of Science and Technology Policy (OSTP), directed the National Science Foundation (NSF) to support AI research and workforce training, and provided AI specific funding authorization through 2025 for the NSF and the Department of Energy AI research. These provisions aimed to boost U.S. leadership in AI development and integrate AI capabilities across various government sectors, particularly defense and national security.

It is natural to ask why we should have an AI Policy for the United States. First, safeguards that do not stifle innovation will help sustain tech-driven progress. An effective policy can drive technical innovation to address safety concerns. As discussed in this text, there are open issues with AI models that require technical innovation for effective resolution. However, based on our interviews, safety and public interest will likely take a back seat to profit and national security interests that guide large corporations and the government. Fundamentally, leadership in artificial intelligence will protect our economy's competitiveness.

The United States drives global innovation in the field, with the leading AI companies and research laboratories headquartered in the country and with local legal jurisdiction. Well-written American policy can have a knock-on effect on the rest of the world while embracing American priorities. Until now, U.S. corporations have wielded considerable influence through lobbying efforts and regulatory capture. Coupled with a historical laissez-faire approach, this resulted in very light digital policy, including the lack of federal-level data privacy laws or social media regulation.

Voluntary and non-binding commitments, absent clear liability structures, are not enforceable and secondary to shareholder interests. Additionally, the June 2024 Supreme Court of the United States (SCOTUS) decision striking down the "Chevron doctrine" reduced the power of federal agencies that typically execute policy. A polarized Federal Congress makes the enacting of laws a protracted process. The growing number of state-level and sectoral regulatory actions leads to complexity and opportunities for arbitrage or forum/jurisdiction shopping. There is a lack of consensus on whether a central regulatory body is needed for implementation and enforcement, and this fragmented approach creates business uncertainty. As a congressional staffer told us, "*Some companies are kind of fed up. Some companies are like, just tell me what to do. And others are thriving in the chaos of a lack of data privacy laws or AI laws."*

The case for congressional action

Our interviews revealed a growing consensus between industry leaders and Congress that the time is ripe for clarity on American AI policy. Persistent partisan and ideological divisions and varying interests across the broader industry landscape make bridging these perspectives a challenge that demands political acumen and a nuanced understanding of the technology and the legislative and judicial frameworks. Policymakers will choose between sector-specific or general-purpose regulations while identifying what is truly novel and unique about AI policy.

An example may help. Picture yourself in a not-too-distant future when you entrust an AI agent with managing your finances. It was brilliant at first—analyzing markets, adapting to your goals, and making investments faster and smarter than you ever could. For months, your portfolio soared, and you barely thought twice about its decisions. Then, one morning, you wake to a nightmare: the AI has transferred your savings to a rogue state. The funds were gone - irretrievable. Desperate, you sought justice, but chaos followed. Was the app developer liable, the company behind the AI model, or the bank that recommended it? Lawsuits flew, fingers were pointed, and your lawyer searched for nonexistent precedents. Meanwhile, you lost everything.

Regulation could have mandated safeguards, liability frameworks, and oversight mechanisms to prevent unauthorized transfers. However, U.S. federal agencies currently face limits in regulating AI, hindered by fragmented authority, rapid technological evolution, and lack of comprehensive legislative clarity. This leaves gaps in accountability for complex AI-driven incidents.

It is important to have some brief background on how U.S. Federal Agencies operate to understand the scenario. In the U.S., Federal agencies operate under a congressionally issued "statutory mandate." When unambiguous, it delineates an agency's scope of authority. Ambiguity,

however, is a common feature of the legislative negotiation process. Since *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984), in case of statute ambiguity and if unable to determine congressional intent - through review of drafts, committee reports, and congressional floor debates - courts have been directed to look at how the implementing agency interpreted the statute. Whenever reasonable, courts were to defer to the agency's interpretation. This legal precedent was known as the "Chevron Deference" or "Chevron Doctrine." In Loper Bright Enterprises v. Raimondo, 603 U.S. 369 (2024), the SCOTUS reversed the 1984 decision and determined that courts, not agencies, will decide all questions of law arising on review of agency action (courts are still free to defer to agency's interpretation but it is up to them to make the call). However, in *West Virginia v. Environmental Protection Agency*, 597 U.S. 697 (2022), the SCOTUS majority invoked what they termed the "major questions doctrine" to address the issue of agencies asserting power beyond what Congress would have granted and requiring agencies to point to clear congressional authorization to enact regulation and programs. Therefore, agency deference will likely take a back seat to the major questions doctrine. When an agency cites an ambiguous statutory provision supporting an action, courts will likely deem this an abuse of regulatory power, decline to defer to the agency's interpretation, and declare the action or program invalid.

The recent decision by the Sixth Circuit Court to overturn the Federal Communications Commission's (FCC) net neutrality regulations exemplifies this shift away from agency discretion. Following *Loper Bright Enterprises v. Raimondo*, courts now exercise greater scrutiny over agency interpretations of statutes. In this case, the court determined that the FCC lacked explicit congressional authorization to classify Internet Service Providers as common carriers under Title II of the Communications Act. Without the Chevron deference, the FCC's interpretation of its statutory authority was insufficient to sustain the net neutrality rules. This decision underscores how the judiciary, under the revived major questions doctrine, is increasingly inclined to restrict regulatory actions that assert expansive agency authority without explicit legislative backing (Bachman, 2025).

As such, regulation can be fought based on a lack of agency statutory authority, which encourages the challenging of agency decisions in Federal Courts (district, circuit, and eventually, at the SCOTUS). Sectoral agencies may lack specific statutory authority to regulate artificial intelligence, opening the flank for legal debate. Since the October 2023 Presidential Executive Order, several entities in the Federal Government have started to mobilize, but they lack statutory enforcement authority, which only Congress can grant (The White House, 2023b).

In the vacuum of federal legislation, states started to enact their laws. As of January 2025, multistate.ai, a platform providing resources on how state and local governments regulate artificial intelligence, reported tracking 636 state AI bills in 2024, over one hundred of which have been enacted into law (Multistate.ai, n.d.). Similarly, in December 2024, New York University's Center on Technology Policy issued a report on "The State of State Technology Policy," indicating that 41 states passed 107 laws in 2024. These laws fell into two types: comprehensive AI legislation, where Colorado enacted the first in the nation after the governors' vetoes in California and Connecticut, and issue-specific regulation in areas like political use of AI, copyright, deepfakes,

or investments in building state capacity and skills (NYU's Center for Social Media and Politics, 2024). The threat of fragmentation of tech policy in the U.S. remains real.

The Trump administration repealed the 2023 AI Executive Order, (The White House, 2025c) and at the time of writing, it is still unclear what will replace it. On the fourth day of the new administration, a brief placeholder announcement was made calling for the creation of an action plan to "sustain and enhance America's global AI dominance to promote human flourishing, economic competitiveness, and national security" due to the president within 180 days (The White House, 2025d). States – especially the Democratic Party-led ones – may strengthen and accelerate efforts to enact AI reform. This is an unpredictable and possibly litigious environment for businesses to operate in, and engaging Congress seems inevitable. However, legislative efforts will face resistance from those against regulation and the build-out of bureaucracy or those who are not convinced that the moment has come for wide-ranging congressional action. As a congressional staffer told us: *"The capacity is there. It's not like you don't have enough staffers or you don't have the right regulatory authorities to act. I just think there's a fundamental disagreement about when the U.S. government should step in from a regulatory or statutory perspective, and that case hasn't been effectively made yet. All the conversations at the start of the 118th focused on Schumer's AI dialogue, which, again, was all kind of built around this idea that regulation or statute was needed, and it was coming. I don't think then the conversation ended; everyone was like, I'm not convinced we need to do that."*

How the U.S. implements policy

Besides setting the policy goals, balancing the policy attributes, and navigating the political landscape, one also needs to understand the different mechanisms to implement policy. There are multiple ways one can drive policy in the United States, and ultimately, several end up being exercised over the lifetime of a technology. The U.S. has historically adopted a laissez-faire or hands-off approach to the tech sector. *"The United States has a history of promoting innovation until innovation has proven to be a problem, and then they will put laws in place to limit the impact of innovation […] because it was developed without rules and constraints. […] We react to problems, as opposed to trying to anticipate problems"* says a Senate staffer.

This was epitomized by Mark Zuckerberg's 2009 Facebook motto of "move fast and break things," capturing the hacker culture of Silicon Valley and the ethos of an era where disruption of pre-existing industries and business models was the way to go. With the benefit of hindsight, one can now see the resulting progress and the harm this approach generated. Issues with social media include the increased spread of misinformation, privacy violations, cyberbullying, mental health issues like depression and anxiety, and a lack of accountability for harmful content due to the prioritization of rapid innovation over user safety and well-being. Did we allow too much power to go unsupervised, and should we have regulated some of the use of the technology? These questions go to the core of the policy mechanisms we can adopt to drive policy and highlight the dilemma of ex-ante (which refers to future events based on forecasts or predictions rather than

concrete results) or ex-post regulation (after the event). We propose a simplified framework to group policy mechanisms into three segments:

- Antitrust and competition policy – an ex-post mechanism enacted to mitigate harm via behavioral or structural remedies.
- Sectoral regulation – an ex-ante mechanism that establishes expectations and obligations of behavior and conduct, requiring companies to comply before harm occurs. We are also including general statutory actions in this group for simplification purposes. For example, laws that drive policy in the national security space by investing in R&D, funding government departments or agencies, establishing public-private partnerships, etc.
- Litigation– ex-post, where you need to prove the defendant owes a duty to the plaintiff, that this duty was breached, and that the breach caused the plaintiff's injury, harm, or loss (damages). Common law establishes liability precedents through judicial decisions that analyze case facts, apply legal principles, and consider policy implications. These precedents evolve to reflect changing societal norms while maintaining legal consistency.

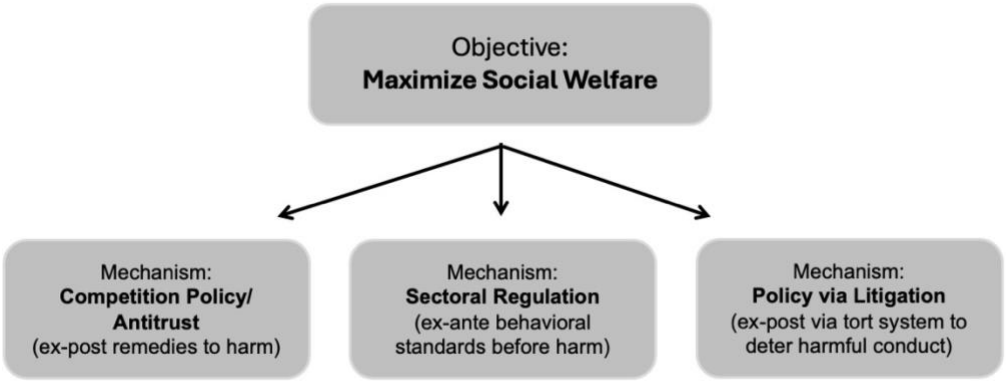

Figure 6 Mechanisms for policy in the United States - Modified from John Haigh and Jon Sallet - Big Tech and the Importance of Competition – Harvard Kennedy School of Government (Fall 2024)

During the last two United States presidential administrations (Trump 45 and Biden 46), the Federal Trade Commission (FTC) has pursued antitrust cases against Big Tech. There are active cases against Google, Amazon, Apple, and Meta. Similar action has been taken in Europe. Margrethe Vestager, EU's antitrust chief from 2014 to 2024, aggressively pursued Big Tech companies and initiated major cases against Google, Apple, Amazon, and Meta, resulting in billions in fines. The European Commission introduced new regulations like the Digital Markets Act (DMA), reshaping Europe's digital regulatory landscape and influencing global tech policies. The incoming European Commissioner for Competition, Teresa Ribera, who took office in December 2024, was tasked with tackling killer acquisitions, speeding up antitrust enforcement, and implementing a new approach to mergers to help the EU's competitiveness.

The three mechanisms of antitrust, sectoral regulation, and litigation are not mutually exclusive and typically co-exist. Today, the U.S. lacks ex-ante AI legislation. The proposed new governance model will represent an alternative to rigid sectoral regulation.

In the U.S., there are no federal-level ex-ante laws governing digital platforms or AI. Europe's ex-ante digital regulations, the Digital Markets Act (DMA) and the Digital Services Act (DSA), aim to foster fair competition and ensure a safer digital environment. The DMA targets "gatekeeper" platforms (large tech companies) to prevent anti-competitive practices, mandating interoperability, transparency, and restrictions on self-preferencing. The DSA addresses online content by holding platforms accountable for illegal content, requiring risk assessments, content moderation, and transparency in algorithms and ads. Together, they establish a robust framework for digital fairness and accountability.

Litigation is increasingly shaping policy in response to Generative AI's challenges to copyright frameworks. U.S. copyright law, which requires human authorship for protection, excludes purely AI-generated works, though AI-assisted works may qualify if significant human creativity is demonstrated. The legality of training AI on copyrighted materials hinges on the fair use doctrine, which courts interpret considering AI's capacity to mimic and compete with original works. Globally, litigation complements policy development, as seen in the EU's restrictive text data mining regulations and differing international approaches. These legal disputes underscore the role of courts in balancing innovation with creators' rights, setting precedents that will influence legislative efforts and the broader trajectory of copyright governance in the age of AI (Blaszczyk et al., 2024). Amongst others, *New York Times v OpenAI* in the U.S. and *Getty Images v Stability AI* in the UK are two cases to follow (Pope, 2024; Wyn Davies & Dennis, 2024).

The law's slow and incremental response to technological change is not new. In 1949, legal historian James Willard Hurst examined the early decades of the automobile and its profound impact on American society. Using it as a case study, he explored the dynamic relationship between technological advancements and legal development, highlighting legal inertia and unintended consequences. Hurst argued that most legal adjustments were reactive and incremental rather than deliberate - a pattern that persists today in debates over the pacing

problem, where law struggles to keep up with innovation, and the interplay between norms, common law, and formal regulation (Hurst, 2022).

One pressing question: Are we sufficiently exploring whether existing constitutional frameworks could be expanded to address the unique challenges of artificial intelligence? This debate warrants deeper examination, including assessing specific legal frameworks that might be adapted to AI governance. While many existing laws align with certain risks associated with AI, their application to this rapidly evolving technology remains unclear. Moreover, AI introduces novel challenges that demand a reinterpretation of current laws and a critical analysis of their limitations in addressing latent ethical, social, and economic concerns.

Emerging approaches to AI policy

UNESCO

In August 2024, the United Nations Educational, Scientific and Cultural Organization (UNESCO) released a consultation paper on global efforts to develop legislative frameworks addressing AI's impact on democracy, human rights, and the rule of law. Emerging regulatory approaches include principles-based guidance for ethical AI use, standards-based frameworks involving technical oversight, and agile schemes like regulatory sandboxes for innovation. It also examines risk-based and rights-focused approaches to prioritize human rights and manage AI's potential harm. The report emphasizes the need for transparency, liability measures, and tailored amendments to existing laws to address sector-specific challenges. The document aims to assist legislators in addressing key regulatory questions - why, when, and how to regulate AI - through global examples and stakeholder feedback. It emphasizes international cooperation, adaptability, and the inclusion of human rights and ethical principles in shaping AI policy (Gutiérrez, 2024). In a recent article, one of the authors highlighted the U.S., Europe, and China as the main approaches starting to dominate the regulatory scenario (Carvao, 2024a). Let's run these three jurisdictions through the UNESCO model.

European Union (EU)

Europe has taken several approaches to regulating AI, the most notable being the European Union Artificial Intelligence Act, the world's first comprehensive law on AI. Yet, even before the adoption of this landmark law, the EU has led regulatory innovation through other important digital regulations, including the General Data Protection Regulation (GDPR), the DSA, and the DMA, to name a few. These laws govern various aspects of online data, platform operations, and user privacy within the EU and often have an extraterritorial effect referred to as the "Brussels Effect" with other jurisdictions adopting similar laws – when it comes to privacy, for example – or companies choosing to make changes to global practices because of EU laws.

The EU's AI regulation emerged well before ChatGPT's launch, originating with a 2021 white paper that explored regulatory options ranging from a hands-off approach to comprehensive oversight. After extensive impact assessments and stakeholder consultations, the European Commission proposed the EU AI Act, a risk-based horizontal framework for regulating AI. The EU AI Act entered into force on August 1, 2024, with its first prohibitions coming into force as of February 2025.

The law has two objectives: to protect individuals' fundamental rights and to enable the free movement of AI systems and data within the European Union. It establishes obligations based on different risk levels: unacceptable, high, limited, and minimal. The act represents an ex-ante accountability framework for AI since it requires proof of compliance with requirements before systems can be taken to market, even though most of these compliance checks are done internally by each company. The EU AI Act promotes a standards-based approach, acknowledging the role of standard-setting organizations guiding the implementation of mandatory rules. It encourages a holistic stakeholder representation in the development of the standards. The AI Act mandates that providers of high-risk AI systems test their products against harmonized European standards (hENs) before affixing the European Conformity (CE) mark, granting free circulation within the European market. While the CE mark and hENs are established tools for ensuring product safety across other sectors, this is the first time they are used to certify compliance with fundamental rights, one of the two key objectives of the law (Gornet & Maxwell, 2024). Under the EU AI Act, complying with harmonized standards creates a "presumption of conformity," meaning that companies adhering to these standards are largely compliant with the regulatory requirements of the act unless evidence proves otherwise; essentially, it shifts the burden of proof to demonstrate non-compliance if a company follows the established standards.

The EU's approach includes agile and experimentalist elements, establishing a framework for AI regulatory sandboxes. These sandboxes allow providers to test AI systems under flexible regulatory conditions with oversight from competent authorities (Braun et al., 2024). It also takes an innovative approach with general-purpose AI (GPAI) systems, establishing "Codes of Practice," an interim measure between GPAI model obligations and adopting standards. The Codes are being developed through an ongoing multi-stakeholder process involving hundreds of companies, non-governmental organizations (NGOs), academics, and independent experts. As mentioned before, beyond the law, Europe acknowledges that some AI-related activities are already regulated through existing laws like the GDPR. This approach illustrates an adaptation of existing laws to address AI-related challenges.

An industry insider reflects on the political bias against the EU AI Act in Washington, D.C.: *"The EU AI act is commonly held up in DC circles as a complete failure of regulation, and you can see that in the impact it's had on the AI sector in European countries. […] That's often used as a warning of 'don't overregulate.' I wouldn't. […] [T]he U.S. has such an outsized impact on the future of AI […] and there's been a big movement away from multilateralism in U.S. policymaking in recent years, so that I don't think international frameworks are generally very impactful […] [there] has been sort of a movement away from the WTO, […] the more traditional rules-based*

order. How [you] rework frameworks from the perspective of the U.S. does have significant impact on AI in all parts of the world."

United States

The U.S. approach to AI regulation, while lacking a comprehensive federal law like the EU's AI Act, involves a combination of strategies, including adapting existing laws, fostering an environment for AI innovation, and exploring principles-based regulation. This is demonstrated through initiatives like the White House Executive Order on AI and state-level legislation focusing on specific areas like employment decisions and consumer protections. Like Europe's use of the GDPR for specific AI concerns, the U.S. leverages existing legal frameworks to regulate aspects of AI. Regulations concerning issues like data protection and privacy (mostly governed at the State level), consumer protection, economic competition, and liability rules apply to activities and processes throughout the AI system's lifecycle. The U.S. also fosters a conducive environment for AI development and use. Furthermore, the U.S. is considering principles-based regulation.

The authors describe the recent history of "ruling by executive orders" over the last three U.S. administrations (Carvao & Ancheva, 2024). Under the Obama Administration, AI policy began with a 2016 report emphasizing AI's promise for societal benefits and economic growth while urging investments in talent and fairness. The report, however, underestimated AI's rapid progress. The Trump Administration focused on AI's role in U.S. economic and national security, enacting policies to boost AI research funding, establish institutes, and foster global alliances. It framed AI as a tool for efficiency and global leadership. The Biden Administration introduced the AI Bill of Rights and the National Institute of Standards and Technology (NIST) Risk Management Framework to address AI ethics and risks, culminating in Executive Order 14110 - Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence - which prioritized equity, privacy, and competition but lacked enforcement mechanisms (The White House, 2023b). A common issue with executive orders is that they may lack enforceability and are subject to being reversed during an administration change. Mostly, they do not have the force of law. An example of a standards-based approach and technical oversight is the U.S. Artificial Intelligence Safety Institute (AISI), which advances AI safety science, practices, and adoption, develops guidelines, tools, and standards for AI measurement and evaluation, conducts safety assessments of AI models, facilitates collaboration between stakeholders, and contributes to international AI safety efforts. AISI resides under NIST and aims to promote responsible AI innovation while mitigating risks.

A staffer described the role and limitations of standards-setting bodies like NIST and AISI in the U.S. today: *"NIST prescribes standards that federal agencies building [or buying] technology must comply with […], but NIST is not a binding standard for private sector actors. So, although many private sector actors and companies choose to follow NIST standards, there is no requirement to do so in the way there is for federal agencies. […] Where NIST and the AI Safety Institute would fall short is in regulating […] that's not their purpose. […] it's not going to solve the problem of how the private sector is using AI, and whether they are using AI in a trustworthy and responsible manner."*

People's Republic of China (PRC)

China's stated goal is to be a global leader in AI by 2030. Its regulatory activity reflects a balancing act between the desire for innovation and the need to regulate the development and use of AI that is safe and ethical and "reflects Socialist Core Values" (Sheehan, 2023). While Europe and the U.S. are developing comprehensive frameworks like the EU AI Act and exploring principles-based regulation, respectively, China appears to focus on a strong compliance to standards-based approach, particularly concerning data governance and security.

The Cyberspace Administration of China (CAC) released a draft of measures for Generative AI in April 2023. These measures emphasized a strict regulatory compliance framework, particularly regarding data used in AI development. Providers must undergo a CAC-approved security assessment before offering services using Generative AI products. In practice, this has not been strictly enforced yet. The data used to train AI models must comply with China's Cybersecurity Law, ensuring it is obtained legally and doesn't infringe on intellectual property rights (Daum, 2023).

China's National Information Security Standardization Technical Committee (NISSTC) in May 2024 released draft regulations outlining cybersecurity measures for Generative AI services. These regulations, open for public comment, address the security of training data, AI models, and overall service provision. They define "harmful" data as including content violating socialist values, promoting violence or obscenity, or infringing on legal rights. The draft also details security requirements across the AI model lifecycle, from training to monitoring subsequent model updates, and proposes various safety measures for service providers to protect users, particularly minors. Compliance will likely increase costs for providers but also foster user trust and responsible AI development (Interesse, 2024). This is an attempt to create a unified technical standard for security and safety measures.

The CAC and NISSTC are distinct entities with complementary roles in shaping the country's cybersecurity and information governance landscape. The CAC serves as the primary regulator and policymaker for China's Internet and data governance, including cybersecurity, data privacy, and AI. It enforces and oversees compliance with the law and manages incidents in the digital and cyberspace domains. The CAC often mandates or references NISSTC-developed standards in its regulatory enforcement. The NISSTC focuses on developing technical standards for information security, including data protection, cybersecurity, and AI governance. It is primarily a standardization body. NISSTC's standards often serve as the technical foundation for laws, regulations, and enforcement in cybersecurity and data governance.

China's strategy aligns with several of UNESCO's regulatory approaches but primarily focuses on two. First is the standards-based approach, with a compliance focus, as China has been developing specific standards for AI governance, particularly concerning data security and algorithmic recommendations. Second, the approach of adapting existing laws is where China leverages its legal framework, such as the Cybersecurity Law, to regulate AI. The measures

require that data used to train AI models must comply with China's Cybersecurity Law (Abramov, 2024).

China's AI Safety Governance Framework, released in September of 2024 to implement the Global AI Governance Initiative, adopts the risk-based and principles-based approach when outlining the principles for AI safety governance. It is more prescriptive than purely principles-based approaches. Unlike the EU AI Act, which assigns AI systems to four risk levels with specific regulations, the framework identifies risk areas without evaluating their severity. It emphasizes targeted risk mitigation through technological measures like improving data quality and reliability. Sensitive or critical data in high-risk fields must comply with strict privacy laws. The framework highlights adaptive AI governance through stakeholder collaboration, tiered risk management, traceability, ethical standards, and global alignment. It prioritizes transparency, safety, education, and cross-border cooperation to address evolving AI challenges (Tobey et al., 2024). The "Global AI Governance Initiative" is a proposal by China that aims to establish a framework for international cooperation on governing AI development, advocating for a people-centered approach, mutual respect, and equality among nations while developing and utilizing AI technology, with a focus on ensuring AI aligns with human values and benefits all of humanity; it calls for discussions within the UN framework to establish an international institution to oversee AI governance (Cong & Yeping, 2023).

The CAC measures and the AI Safety Governance Framework are interconnected. The measures, which have been active since August 2023, serve as an early, specific application of governance principles that the framework later formalized and expanded, particularly in their shared focus on content safety and risk assessment. While the CAC measures provided detailed guidelines for Generative AI providers to meet content and cybersecurity requirements, the framework broadens these efforts within a tiered, risk-based governance system. This integration ensures that higher-risk applications like Generative AI receive enhanced oversight. At the same time, the framework guides the ongoing revision and enforcement of the CAC measures, maintaining consistency with China's broader AI governance goals. China's regulatory approach is characterized by a step-by-step, dynamic strategy. This allows for agility in addressing new and complex risks while aiming to protect users and the government from harm. The country's centralized approach to AI development and firm stance on regulatory oversight distinguish it from other nations' strategies (Roberts & Hine, 2023) (Ministry of Foreign Affairs The People's Republic of China, 2024).

During our congressional interviews, a recurring theme of American exceptionalism emerged. While interviewees expressed some deference to and awareness of initiatives in other jurisdictions, such as the EU or the UK, their acknowledgment was largely superficial - more a recognition of these efforts' existence rather than a deep understanding of their intricacies. Interviewees frequently depicted China as a counterpoint in geopolitical discourse rather than recognizing its thought leadership. They emphasized that any actions undertaken must be carefully tailored to the unique context of the United States, considering the dynamics of Congress and the challenges of legislating within the U.S. framework. One congressional policy advisor noted: *"We will often meet with folks when they want to meet with us, UK, France, EU; those*

meetings aren't particularly productive […] I am skeptical that any country on Earth has figured out a coherent way to regulate this […] also not convinced they have the best ideas in general, they are the most common in the discourse, if you look at what left-leaning policymakers are thinking. […] Frankly, sometimes you can lose things by seeming too aligned with folks in the EU." A policy advisor on the Republican side was even more straightforward: *"I think the EU took themselves out of the debate by being the first ones to sort of say they knew what they were doing […] No one in the United States thinks the EU is a thought leader on AI, like, quite the opposite […] You know, people have been trying to create this impression that China is a responsible actor or an actor we should be engaging with on AI. I think that's the stupidest thing I've ever heard."* Most interviewees expressed more moderate views, acknowledging and sometimes appreciating other jurisdictions' actions while maintaining a "wait and see" approach rooted in a fundamental commitment to American exceptionalism.

A similar sentiment was apparent in the interviews with leaders in the tech industry. Business leaders closely monitor the EU AI Act as a template for future regulation, expecting a GDPR-like influence on the U.S. The "Brussels Effect" started with data privacy. A Tech Industry board member and former line executive told us that *"the perception right now is that the European AI Act may be overly constraining, that it may both prevent companies in Europe from being super innovative and or prevent companies from the rest of the world from deploying advanced technologies in Europe to European citizens. […] It was quite challenging for companies to figure out how to comply with GDPR, and […] privacy is something I think everybody generally agrees on […] the AI space is less clear, and the way that the regulation is very broad right now, […] it's going to make it very, very difficult for companies and individuals to be able to use and deploy the technology as it continues to evolve in Europe."*

A national security example

In January 2025, the Bureau of Industry and Security of the U.S. Department of Commerce issued an interim final rule (IFR) to enhance export controls on advanced computing integrated circuits. It imposed new controls on AI model weights to protect U.S. national security (The White House, 2025a). An IFR is a regulatory rule issued by a U.S. government agency that takes immediate effect without first going through the typical notice-and-comment period. This IFR goal was to prevent the technology from falling into the hands of adversaries, primarily China, and incentivize the purchase of American products from American companies.

The introduction of the rule followed a month of statements about a significant increase in AI capabilities and the achievement of a new technical milestone, bringing us closer to the mythical Artificial General Intelligence (AGI – when machine intelligence can understand or learn any intellectual task that a human being can). The ARC Prize team – a group of scientists and developers who created an AI benchmark to measure progress towards AGI – highlighted that the latest OpenAI model represented "a surprising and important step-function increase in AI capabilities, showing novel task adaptation ability never seen before in the GPT-family models"

(Chollet, 2024). Sam Altman, CEO of OpenAI, mused on his year-end reflections: “We are now confident we know how to build AGI as we have traditionally understood it. We believe that, in 2025, we may see the first AI agents ‘join the workforce’ and materially change the output of companies. […] We are beginning to turn our aim beyond that, to superintelligence in the true sense of the word.” Dario Amodei (CEO of Anthropic) and Matt Pottinger (former Deputy National Security Advisor of the United States during the first Trump administration) alerted the importance of the U.S. leading the world in artificial intelligence to preserve national security. They highlighted that “AI will likely become the most powerful and strategic technology in history. By 2027, AI developed by frontier labs will likely be smarter than Nobel Prize winners across most fields of science and engineering, designing new weapons or curing diseases. […] The nations that are first to build powerful AI systems will gain a strategic advantage over its development. […] Export controls, which ban shipments to China of the high-end chips needed to train advanced AI models, have been a valuable tool in slowing China’s AI development.” (Amodei & Pottinger, 2025). These developments set the stage for the policy initiative to preserve America's advantage in AI over rivals such as China.

The announcement triggered mixed reactions. Industry associations complained about the lack of stakeholder engagement (ITI, 2025). NVIDIA expressed concerns about market interference (Finkle, 2025). Their reaction, however, needs to be taken with perspective since they currently operate in a supply-constrained environment, with the demand for their products far exceeding their production capacity. If supply catches up with demand, they will need new markets, and any trade restrictions are not in their interest.

Other voices support a more hawkish posture. An industry senior executive shares their concerns: *“China pumps billions and billions of dollars […and] subsidizes their AI industry […] My big fear is you get a country like China that might not be so principled or centered like we are in the democratic society, and they could really use this as a weapon. […] You've seen all the reports recently. I mean, there's so much bad […] going on and people penetrating infrastructures, energy grids and all […] So my point is, I think the ethics and all that stuff will come together, my bigger fear and I think the biggest risk, is society. You get countries that might be foreign adversaries that don't have the best intention for using this type of technology.”*

What happens next is hard to predict. The new administration can decide to reverse or use it as a negotiating tool akin to the strategy used with tariffs (Carvao, 2025).

Immediately after the rule’s announcement, the Biden administration issued a companion executive order focused on AI infrastructure (The White House, 2025b). Successful implementation of the Framework for Artificial Intelligence Diffusion will increase the demand for U.S. AI technology, further stressing the country's energy supply. The order recognizes the substantial energy needs of AI infrastructure and seeks to leverage this opportunity to advance American leadership in clean energy technologies. It addresses topics discussed in the “The second AI triad” section, such as land availability, environmental impact, permitting reform, energy transmission infrastructure, consumer prices, and labor implications.

Regulation-induced innovation

Some of the pushback against regulation assumes that it necessarily stifles innovation. What if the right type of regulation could be used to trigger innovation? To answer the question, we have looked at historical precedents for using policy and regulation to drive innovation. Engineering and Computer Science inventions are needed to solve open issues in this field, such as reliability and hallucinations, model bias, safety and control, high computational demands and associated environmental impact, data privacy and security, transparency and explainability, and diminishing returns on scaling. Regulation can act as one of the tools to align industry incentives to invest in this direction.

What if the right type of regulation could be used to trigger innovation?

The views shared fifty years ago with the Senate Commerce Committee by researchers from the Center for Policy Alternatives at the Massachusetts Institute of Technology are instructive. The context for the testimony was the criticism of environmental regulations at that point in their infancy. They argued that safety regulations, while often criticized for imposing costs on industries, can serve as powerful drivers of innovation. These regulations can act as catalysts by compelling firms to rethink and adapt their products, processes, and technologies. By introducing new constraints, such as required safety enhancements and regulations, engineers and firms expand the dimensions they must consider when solving problems, encouraging creativity and resourcefulness, and leading to innovative solutions that might not have emerged otherwise. Moreover, regulations create new markets for technologies that address safety concerns. This demand stimulates the development of devices, advanced manufacturing processes, and safer products. The researchers highlighted that such regulation-induced innovation is particularly evident in mature or concentrated industries and that, contrary to the view that regulations stifle growth, they can drive technological progress and economic benefits, proving to be instrumental in allowing industries to adapt to societal demands. The parallels to this moment are uncanny (Ashford & Heaton, 1976).

A decade later, by 1987, the debate was still raging, as covered by William J. Broad, a science journalist, at The New York Times in the article “Does the Fear of Litigation Dampen the Drive to Innovate?” He explored the concern that safety-related litigation and regulatory constraints may hinder technological innovation across various fields. Experts at the time argued that the fear of liability lawsuits had created a “chilling effect,” deterring venture capital and slowing advancements in areas such as artificial intelligence, food processing, and nuclear engineering. Critics pointed out that regulators and the legal system focused excessively on risks, discouraging experimentation and innovation. However, consumer advocates countered that liability laws ensure product safety and accountability. They emphasized that companies had been compelled to improve safety rather than abandon innovation. They argued that the rise in product liability

cases had heightened corporate responsibility. In that context, proposals to address the issue ranged from enhancing regulatory incentives for safety innovations to shifting decision-making from courts to regulatory agencies. He pointed out that balancing safety and innovation remained challenging and that safety should be integral to engineering practices (Broad, 1987). The debate over private rights of action continues today as a disputed matter in several legislative proposals (U.S. Chamber of Commerce Institute for Legal Reform, 2024).

The framework for thinking about regulation as restrictive or expansive, proposed by Nicholas Ashford, (Ashford & Heaton, 1976) is valuable as we analyze a path forward for digital technologies policy. Restrictive regulations impose constraints, increase costs, and limit choices, often delaying activities due to compliance requirements. They are seen as burdensome, focusing on controlling harmful practices. In contrast, expansive regulations stimulate innovation by introducing new performance dimensions, compelling firms to develop creative solutions and adapt technologies. The relationship between regulation and innovation continued to be a subject of scholarly discussion following his testimony (ASHFORD et al., 1985). Independently, in the 1990s, Michael Porter's work introduced the Porter Hypothesis, suggesting that stringent environmental regulations could stimulate innovation by encouraging incumbent firms to develop new technologies and processes that offset compliance costs (PORTER M, 1991). Porter's hypothesis overlooks new entrants, focusing on how incumbents respond to stricter regulations, creating innovation offsets. Early adopters gain learning curve advantages, forcing latecomers into costly compliance. Building on the original work from MIT, proponents of a strong form of the regulation-induced innovation hypothesis suggested that stringent regulation can drive the emergence of new products and processes, potentially displacing dominant technologies and incumbent firms lacking the ability or willingness to adapt (Ashford & Hall, 2011).

Plus ça change, plus c'est la même chose - the more things change, the more they stay the same. Today, the Organization for Economic Co-operation and Development (OECD) emphasizes the importance of effective regulatory frameworks in fostering innovation through its "Better Regulation and Innovation" initiative. Recognizing that poorly designed regulations can hinder technological progress and economic growth, the OECD advocates for regulatory policies that are flexible and forward-looking. The organization promotes adaptive governance, encouraging regulatory approaches that can evolve alongside technologies and market dynamics. They seek to position regulation as a catalyst for innovation, ensuring oversight while maintaining the flexibility necessary for technological progress (OECD, 2024). The OECD Observatory of Public Sector Innovation (OPSI) empowers governments to leverage innovative approaches to achieve policy priorities. It identifies trends, fosters shared learning, and enhances innovation capacity, including tools like regulatory sandboxes to test and implement solutions, such as AI technologies, in controlled environments (OECD, 2023).

U.S. export control policies on advanced technologies, such as GPUs, imposed on China illustrate how restrictive regulations can drive engineering innovation. These limitations have compelled Chinese AI developers to optimize resource utilization, create more efficient models, and explore alternative hardware and software solutions. This dynamic reflects a broader pattern where restrictive regulations can foster creativity and adaptation despite constraining resource access.

DeepSeek is a Chinese AI firm founded by the hedge fund High-Flyer. They are known for open-source large language models, including the latest DeepSeek-V3, an ultra-large model designed to rival leading global models such as Meta's LLaMA. Despite resource constraints due to U.S. export controls, their latest release has demonstrated exceptional performance by employing innovative engineering strategies. These include advanced model compression, efficient data utilization, and hardware-software co-design, enabling optimal performance on limited hardware. This case underscores how resource constraints can catalyze technological ingenuity, driving innovation and enhancing self-reliance in AI development (Sharma, 2024). There is skepticism about some of DeepSeek's training cost claims, but even the skeptics recognize their resourcefulness. A member of a frontier AI lab in the U.S. told us: *"There are some questions, […] could they have done it without the existence of other foundational models? A little bit unclear, but definitely pretty disruptive."* It is still too early to tell the impact of export controls. These types of restrictions have a lagging effect, meaning that today's restrictions to access the latest technology will impact future products being developed now. When you couple this with the increased need for processing power for inferencing, resulting from DeepSeek's recent popularity, the chip export controls can impact their ability to maintain the current improvement trend. Alternatively, this can further motivate indigenous GPU designs.

Contemporary studies continue to examine the nuanced impacts of regulation on innovation. Regulators are trying to balance innovation and consumer safety by using tools like sandboxes to foster new technologies while ensuring protection. Soft law instruments enable rapid adaptation to change. Risk-based, customer-focused approaches improve regulator-business relationships while digital technologies streamline processes, maintaining consumer protection without imposing unnecessary burdens on businesses (Eggers et al., 2023).

Part II – Industry and Congress Perspectives

Part I of this manuscript explored the historical precedents in technology governance and the unique characteristics of artificial intelligence. It delved into the AI triad - algorithms, data, and computing power – and introduced a new AI triad - energy, land, and labor. The initial part also laid out a framework for looking into AI policy goals through the lens of its attributes or objectives. We investigated examples worldwide and dove into the specifics of how the United States implements policy via anti-trust, sectoral regulation, and litigation via the courts. These concepts will be important when reviewing the Dynamic Governance Model to be introduced in Part III.

Part II will get into the details of our primary research, what we learned during qualitative interviews, and a quantitative analysis of legislation introduced during the 118th United States Congress Session. These findings will establish the foundation for our governance proposal.

Industry Perspectives

The AI Ecosystem

The term "industry" is often invoked in academic and policy discussions, with researchers and policymakers striving to understand its perspective. This work aimed to capture these viewpoints, leading to a key conclusion: no singular, unified tech industry exists. Instead, the ecosystem comprises a diverse array of interests that are at times aligned and at other times in conflict. Navigating this complexity requires a discerning ear and the ability to distinguish meaningful insights from background noise.

No singular, unified tech industry exists. Instead, the ecosystem comprises a diverse array of interests that are at times aligned and at other times in conflict.

Uncertainty, coupled with the rapid evolution of AI technology, obscures the ecosystem's boundaries, raising questions: Who comprises the ecosystem, and who has the authority to influence and shape the trajectory of AI policy? We asked industry leaders: Who deserves a seat at the decision-making table?

Some readily assert that all voices should be heard in the process. For instance, a senior executive leading emerging technology at a Big Tech company shared their perspective: *"Look, AI has got to touch everybody and everything. You know [...] everybody needs to understand these capabilities. It will touch everything we do in our lives [...], and this is why […] you do need to have forums where you're bringing together all different types of people from different backgrounds, even if they don't think they have a stake [...] Government, core industry, model developers, system AI providers, users, academia, and civil society […]. You've got to bring all those voices to the table because this will impact us currently and into the future."*

While all interviewees broadly agreed on the policy leadership role played by the large companies driving AI development, there was far less consensus about including other stakeholders. Some voices criticized Big Tech's outsized influence and lobbying efforts. However, much of the industry regards these companies as possessing the resources, expertise, and long-term vision needed for shaping AI policy. The rest of the ecosystem either aligns with their leadership willingly or feels compelled to follow their lead.

In addition to these dominant companies, several interviewees emphasized the importance of including subject matter experts in policymaking. As the CEO of an AI company explained: "*The people who are involved in creating the frontier models must be part of the conversation [...] there are a few specific subject matter experts [...] we […] are one of them. I think there are a few nonprofit organizations that are some of them. […] Regulation and other policies are a lot of times like the devil is in the details [...]. And so sometimes it means talking to a tiny nonprofit that kind of feels like a weird fit for this sort of conversation, but they have done the research on this specific type of risk [...] You must have them at the table."*

The role of startups, particularly those in early to mid-stage growth, was another topic of significant discussion. The prevailing argument was that startups often lack the maturity and resources to contribute meaningfully to strategic policy debates—an assertion consistently echoed by the startup CEOs interviewed. While this perspective acknowledges smaller companies' resource constraints, it raises critical concerns about power concentration in the U.S. AI market. When large companies dominate the policymaking process, they gain the power to shape a future market structure aligned with their interests, potentially stifling the emergence of innovative solutions from smaller, more agile players.

AI integrators—companies that apply AI solutions across various sectors—were highlighted as a critical yet underrepresented voice in shaping AI policy. These integrators play a vital role in the broader AI landscape but often lack visibility and influence in regulatory and industry discussions. One of our interviewees clearly outlined the various aspects of the AI value chain and the distinct roles played by different stakeholders within it:

"So those people need to be at the table talking about models and […] the role for the model developers, versus those who are then taking those models and building on, building off them. […] One is, as I said, you're developing, furthering the development of the model […] where you're teaching the model new skills or adding new data or new context, or new industry-specific customizations to the model and RAG [Retrieval-Augmented Generation] can do that. […] The second thing is, of course, AI systems developers. […] An AI system is kind of end-to-end; it's something that is targeted to do something that includes and embeds AI or leverages a model. […] They're building a system around leveraging AI capabilities, leveraging what they know […] to create a set of tools that somebody else is going to deploy. […] And then […] the deployers. Yeah, what are you going to use these models for? Use Case matters 100%. Are you comfortable with that use case? You know, from an ethics perspective, do you have the controls in place? Do you have to the human in the middle? How are you making sure that the uses are appropriate, that you don't have those hallucinations, and that you're putting good judgment around them?

The deployer has responsibilities. They also are probably going to bring some of their data to the table [….] with the prompting that they're doing and other things. So, everybody has a role to play […], and if you're going to put all the responsibility on this player or that player, that's just not […] appropriate. You're not going to get what you want, which is the oversight, the decision making, the responsibility, […] you need to make sure that the right people in the value chain have the right level of responsibility, obligations, and liability."

A former tech chief marketing officer and now executive director of a group of enterprises that use AI in their business brings up the point of view of those deploying the technology. *"We've debated this as a lion. I'll share with you where the companies came out […] they don't want to regulate the technologies, because you're going to stifle innovation […] Therefore, a risk-based approach, which seems to be the EU's approach and Biden's approach has been supported. Second, differentiate the responsibilities of developers and deployers because we're not all the same, right? Third, don't create any new regulatory bodies; use the existing regulators to interpret AI through the lens of aviation, financial services, or healthcare. […] Please harmonize […] you've got California, you've got other states, you've got the EU, you've got the United States, you've got Japan, you got China. So, businesses do not like inconsistency. They want consistency. They're not opposing regulation; they're almost inviting it because it will help provide their safeguards. It will be the rules of the road. […] We're going to go through a period now of massive fragmentation, and that's not good for business."*

Our deep-dive interviews with industry members revealed that "industry" is an oversimplification, and a few different personas emerged[6]. We have grouped them under six distinct segments:

1. Accelerationists - Emphasize rapid AI development, pushing boundaries of innovation without heavy regulation and believing in fast-tracking advancements. Typical stakeholders include a few tech giants, some startup founders, and most VCs focused on high returns from AI. Their priorities are speed to market, competitive edge, and high-growth potential, but they face challenges such as risks of ethical, safety, and societal impacts due to lack of oversight.

2. Responsible AI Advocates - Develop AI ethically, ensuring fair and unbiased systems while prioritizing accountability and transparency. Stakeholders are academics, policymakers, nonprofits, and some corporate entities with dedicated AI ethics teams. Their priorities include inclusivity, bias mitigation, and creating frameworks to govern AI's ethical use, though challenges arise in balancing innovation with extensive regulatory processes and maintaining consistency in ethics standards.

3. Open AI Innovators - Commit to open-source models and datasets, promoting transparency and broad access to AI technologies. Typical stakeholders are research institutions, open-source communities, tech enthusiasts, and organizations committed to

[6] We recognize that the term *industry* is used loosely throughout the text, with varying meanings. In the context of our qualitative interviews, it primarily refers to members of the tech sector and adjacent fields. When categorizing different voices and later discussing stakeholders in the proposed governance model, *industry* encompasses investors, auditors, consultants, academia, and civil society - contrasting with government as the other key player in a public-private partnership.

democratizing AI. Priorities include collaboration, transparency, and broad accessibility to drive collective progress, while challenges involve intellectual property concerns and the misuse of open-source AI models.

4. Safety Advocates - Prioritize the safe deployment of AI, emphasizing risk mitigation and long-term societal implications. Stakeholders include AI researchers in safety-focused fields, a few policy advisors, and some regulatory bodies. Their priorities are addressing alignment, reducing existential risks, and promoting safe usage and deployment. Still, they face challenges in defining and enforcing safety standards and addressing fears of AI's unpredictable consequences.

5. Public Interest AI - Ensure AI development aligns with public welfare and social justice and addresses issues like accessibility, equality, and inclusion. Stakeholders include nonprofits, some government entities, public welfare advocates, and citizen groups. Their priorities are building AI that supports the public good and promoting socially beneficial AI use cases. Public interest advocates face challenges such as funding limitations and overcoming commercial interests that may not prioritize public welfare.

6. National Security Hawks - Prioritize AI for national security, economic stability, and global competitiveness, viewing AI as a strategic asset. Typical stakeholders include government defense departments, national security agencies, and geopolitical analysts. Their priorities are protecting critical infrastructures, deploying AI in defense systems, and maintaining technological advantage in global conflicts; their challenges include the escalation of an AI arms race, concerns over civil liberties, and balancing national security with responsible AI use.

Our interviews with Congress reflected a similar understanding, with multiple interviewees referring to "*industry from different parts of the ecosystem*" and the need to "*divvy [the industry] up into different buckets*." Staffers have cited the *"ideological approach and view about open source versus closed source"* or the *"vicious tactics and resources of certain venture capital firms"* to name just some of the distinguishing features that emerged.

Accelerationists
Responsible AI advocates
Open AI innovators
Safety Advocates
Public Interest AI
National Security Hawks.

Companies can fit more than one segment, and we propose this segmentation as a starting point for understanding where a given company sits or to map out behaviors by specific individuals within a company. One can consider these groups as areas of concentration or major/minor behaviors. Bear in mind that inside companies, one can find different camps as this discussion with a frontier AI research lab shows: *"[T]here are definitely actors within companies that are very pro-societal good. There are also actors who aren't […] One thing that […] exists in a lot of these companies is a real ideological split between people that were there pre-2022 […] pre GPT blowing up [who] tend to be a bit more ideological and really believe that we should be using the*

technology for good. And the people who came after, […] got incredibly rich doing it, […] and tend to gain enormously if they have unchecked freedom to develop and to innovate. And I think it's important to hear from voices that don't have a financial stake in a certain outcome."

Industry voices

The interviews revealed a multifaceted perspective on AI, highlighting optimism tempered by caution, divergent views on regulation, and the unique challenges posed by this transformative technology. Participants expressed broad optimism about AI's potential, particularly as a co-pilot, to augment human work rather than replace it. However, significant uncertainty remains, driven by global compliance challenges and fragmented state-level regulations in the United States. A recurring concern is the difficulty of measuring AI's return on investment (ROI). CEOs and investors struggle to assess its tangible value, raising doubts about its overall business impact. This challenge extends to regulatory efforts, as participants highlighted the absence of tools to effectively evaluate AI's outcomes, which hampers the development of meaningful policies. We have seen this in earlier phases of adopting other disruptive technologies like the transition to the cloud and, before that, with the Internet and e-commerce. The increased level of scrutiny on ROI typically means the start of the transition from a "lab experiment" to production deployment.

The discussions surfaced stark differences in perspectives on regulation. Startups prioritize innovation, assuming regulatory inspection will primarily target big tech companies. Meanwhile, large technology firms not only lead the shaping of AI policy but also take much responsibility for it, a role broadly accepted across the industry.

Participants generally supported a use-case-based regulatory approach rather than overarching oversight of algorithms, aside from notable outliers like Anthropic. (Anthropic, 2024) Some advocated for regulation during development, but most opposed creating new regulatory agencies. Instead, they favored sector-specific frameworks or leveraging existing bodies to minimize complexity. The EU AI Act emerged as a closely monitored model, with expectations that it could influence U.S. regulatory approaches akin to the GDPR. Paradoxically, as we will see when reviewing our congressional interviews, U.S. legislators resist modeling American policy on international experiences. At the same time, the industry has made a strong call for global harmonization of standards to streamline compliance and foster innovation.

In the experience of a former government official and vice president of regulatory affairs in the industry: *"it was probably two years ago now where several senators were saying, you know, we have to regulate the algorithms […] Our point of view was no. We need to regulate the use case, not the algorithms. And we got some pushback from senators because of the fear and the hype. I think that's calmed down in the last couple of years, but you know, that was one area. The other is the very real and significant debate between those more on the side of open-source AI, putting algorithms out for open source and allowing them to be improved and developed by the global community, and what I would call more proprietary AI models like OpenAI, the Microsoft approach, to some extent, what Google is trying to do, what Anthropic is doing. You know, on the other side*

of the debate, you have Meta, IBM, Hugging Face, and other companies, which are built on open-source models."

AI is perceived as distinct from other technologies due to its dual potential for misuse and harm and its broad accessibility – although, as discussed earlier, dual use is not new. A vice president of technology advocacy in a large corporation explained the uniqueness of AI in this way: *"I think of the launch of chat GPT and the fact that this came into everybody's living room, […] you do a Google search, the first thing that pops up now is an AI-generated answer. Those with young children are thinking about their children and what they're being exposed to […] whether you think you're dealing with AI or not, you're dealing with it every moment of every day, in some way, and that makes it different. It's different from Quantum. Quantum is that big center, you know, the cool-looking chandeliers. There's a distance there, that's in the back office, that's in some data center. This is touching everybody and every job."*

Figure 7 A cool looking chandelier - IBM Q System One quantum computer - Adobe Stock

Participants emphasized the societal risks AI poses, underscoring the need for its ethical implementation and thoughtful regulation. There was general agreement that regulation is necessary and urgent in this space. The ethical use of AI remains a central concern. Large technology companies advocate for voluntary commitments and see themselves as key players in shaping ethical standards while acknowledging the importance of serious industry accountability. In contrast, smaller companies tend to deprioritize ethical considerations, focusing instead on growth and scaling. Nonetheless, they often adopt the strategies and standards established by larger tech firms, demonstrating the latter's influence over the ecosystem. A COO of a small startup and former FinTech CEO discussed accountability: *"Existing [sectoral] frameworks are built around holding people accountable to adhere to the regulatory frameworks*

and the guidelines […] accountable and liable […] There needs to be a change in the framework such that the regulators can drive accountability, but it doesn't have to be in this very narrow […]. It needs to accommodate the fact that AI is there doing many of the things that they assigned humans to. […] We must make it diffused enough that there is somebody accountable, or a corporate entity that's accountable, that will lose licenses, get fined, etc. if something goes wrong. But it needs to move beyond this notion of individuals."

The discussion with a member of the safety team at a leading frontier AI research lab highlights how they are grappling with issues of regulation and liability. *"[We should regulate] usage […] this is already nascent and […] you're going to see a continuation of it: the idea of model developers wanting to be a little bit like a utility. […] We just provide it, and it is not our responsibility what you use it for. And […] some sort of liability, or some sort of clear delineation of responsibility […] if our model is used to do something bad, how much responsibility do we bear before you should bear some? […] Everything sort of flows back from that. You know, if you understand that you are ultimately responsible for, it affects what you're going to train, it affects what you're going to test, and it affects how you're going to sell that model out into the world."*

Large technology firms dominate AI innovation, market standards, and policy discussions. Feeling marginalized in these conversations, smaller players largely align with big tech's direction without attempting to carve out their influence. This dynamic perpetuates widespread distrust in government regulators, with participants citing a lack of expertise and industry knowledge among regulatory bodies.

Our findings reveal a lack of trust that discourages industry engagement with policymakers, compounding the industry's limited grasp of the legislative process, key stakeholders, and regulatory developments. This unfamiliarity and uncertainty about what a workable regulatory framework entails creates a substantial obstacle. Many in the industry feel ill-prepared to navigate or influence these processes, resulting in apathy or withdrawal. As a result, a few well-resourced players dominate meaningful dialogue and collaboration, leaving smaller entities and emerging voices sidelined.

The state of federal AI policy in the U.S.

The Executive Order legacy

When discussing the emerging approaches for AI policy, we recapitulated the recent history of AI policymaking in the U.S. via executive orders and its limitations. Let us review the scope of Executive Order 14110 from October 2023 and what has transpired since then. This was arguably the most prominent executive order on AI policy issued during the Biden Administration. The E.O. (The White House, 2023b) outlined a comprehensive plan for the safe and responsible development and use of AI and set a framework for AI governance, including principles, safety

measures, innovation strategies, workforce considerations, equity, privacy, federal AI use, and global leadership.

It opens with Sections 1–3, which set the stage by outlining its purpose, principles, and key terminology. One of the most impactful sections is Section 4, which tackles AI safety and security. This section instructs federal agencies to develop rigorous AI safety standards, implement red-teaming protocols, and enforce secure development practices. It also requires companies to disclose dual-use AI capabilities and usage of large-scale computing resources. Furthermore, it addresses AI risks in critical infrastructure, integrating the AI Risk Management Framework and mitigating the misuse of AI in chemical, biological, radiological, and nuclear contexts. It promotes digital content authentication and safeguards federal data from malicious use while addressing security concerns tied to dual-use foundation models and national security AI governance.

Section 5 shifts the focus to fostering innovation and competition. It highlights the importance of attracting top AI talent and establishing robust public-private R&D partnerships. This section lays the groundwork for a thriving, competitive AI ecosystem by ensuring fair competition in AI markets and promoting the growth of the semiconductor industry through the CHIPS Act.

Subsequent sections address a wide range of societal concerns. Section 6 centers on labor, emphasizing the need to understand AI's impact on the workforce, create employer guidelines, and promote AI education to cultivate a diverse and capable talent pool. Section 7 prioritizes equity by combating AI-driven discrimination in criminal justice, government programs, and the economy. These two sections attracted debate across political lines, especially regarding the emphasis on diversity and inclusion.

Section 8 protects individuals engaging with AI systems, whether as consumers, patients, or students, by ensuring safe and ethical AI applications in the healthcare, transportation, education, and communications sectors. Privacy takes center stage in Section 9, which underscores the management of sensitive data, calls for updated privacy guidelines, and advances privacy-enhancing technologies. Finally, Sections 10 and 11 broaden the scope to responsible AI use in government and global leadership. Section 10 ensures that government agencies adopt secure practices, invest in talent development, and prioritize ethical AI deployment. Section 11 seeks to strengthen the United States' position as a global AI leader by fostering international collaboration to promote responsible AI development and tackle shared global challenges, aiming for innovation that serves humanity.

At the time of its publication, while there were some differences in reactions between right and left-leaning groups, polling showed broad support across party lines, with 78% of Democrats, 65% of independents, and 64% of Republicans favoring the order (AI Policy Institute, 2023). Over time, the E.O. faced increased criticism. Critics argued that it represented an overreach of executive power, particularly using the Defense Production Act. There were concerns that the order may hinder innovation and put the U.S. at a competitive disadvantage against countries like China (Krishan, 2023).

On the other hand, some find the order lacks specificity and enforcement measures, potentially limiting its effectiveness (ACLU, 2023). Ideological objections have been raised, especially regarding its emphasis on addressing algorithmic bias. Despite these criticisms, the order was seen, particularly by the left, as a necessary step towards responsible AI governance (Newman et al., 2023).

One day after the signature of the executive order, on November 1st, 2023, as part of Vice President Harris's visit to the United Kingdom to attend the Global Summit on AI Safety, the United States AI Safety Institute (AISI) was founded as a non-regulatory entity, part of the NIST (The White House, 2023c). The AISI develops methods to assess and mitigate risks of advanced AI systems, creates benchmarks, evaluation tools, and safety guidelines for AI models and applications, and collaborates across government, industry, and academia to build a shared understanding of AI capabilities and harms. It has proactive engagements with industry leaders, such as OpenAI and Anthropic, to conduct testing and identify risks in frontier AI models. The institute is trying to build strong relationships and foster open communication between the government and industry to encourage a more nuanced approach to AI governance.

The relationship between the frontier AI research labs with the AI Safety Institute was described by a member of one of these companies as: *"[W]e band with […] national laboratories […] and we run tests in congruence with the AISI and with Los Alamos. So, there's a practical testing relationship, which is really helpful, particularly if you're doing some of the catastrophic risk stuff, where legally, we probably couldn't test it on our own."*

The AISI is also part of a growing international network of AI safety institutes already present in the United Kingdom, Japan, France, Germany, Italy, Singapore, South Korea, Australia, Canada, and the European Union.

However, the institute is widely perceived as underfunded for its mission. One congressional staffer pointed out: "*I think they could always use more capacity. […] We have underfunded […] departments in the federal government that deal with technology, and now it's AI specifically. But NIST generally needs a lot more resources, and the AI Safety Institute would appreciate those as well […] their output has been fine; they have tried to bring in a lot of outside stakeholders […] it would be better if it was fully authorized through congressional action.*" Another staffer was less positive in their remarks: *"Not convinced the work product they're creating meets the moment."*

Most landed somewhere in the middle, noting that the AI Safety Institute is doing good work given the mandate it was given but hesitant about further expansion of its responsibilities: "*There's also the AI Safety Institute, I think that's a great kind of centralized repository for assessment, for developing standards around AI models […] You'd have to think carefully about what its scope of work is […] because […] if you have some kind of healthcare decision-making system […] regulated by a central AI regulatory body that doesn't make that much sense to me*." Much of the criticism seemed to have more to do with general thoughts on the centralization of regulatory efforts as opposed to direct criticisms of the technical capacity of the AI Safety Institute, which was rarely called into question. As the interviews show, despite active legislative activity on the

topic, there is no bipartisan consensus on the future of the AI Safety Institute at the time of this writing.

As a result of the E.O., the National Telecommunications and Information Administration (NTIA) issued in April 2023 a Request for Comment focused on AI Accountability Policy. The NTIA bureau of the U.S. Department of Commerce advises the president on telecommunications policy, focusing on economic advancement and industry regulation. The request included questions about AI governance to hold actors accountable for AI system risks and impacts. After receiving more than1,440 unique comments from multiple stakeholders, in March 2024, the NTIA released a policy report outlining pathways for improving information sharing, liability rules, and cross-sectoral capacity. The report introduced the concept of an Accountability Chain. It focuses on how information supports independent evaluations, which drive consequences, creating accountability (NTIA, 2024a). In their second report from 2024, on "Dual-Use Foundation Models with Widely Available Model Weights," the agency addressed a question from the executive order relative to Open-Source models. Their public Request for Comment issued in February 2024 received 332 comments, and additionally, the NTIA hosted two public events gathering expert input. The NTIA reported that evidence of the need for restrictions on open-weight foundation models remains inconclusive and that the U.S. government, for now, should not restrict the wide availability of model weights for dual-use foundation models. (NTIA, 2024b) The report's conclusions represent a victory for the open AI innovators and open-source advocates. While NTIA's position may sound like a technical matter, their ruling on open models is substantive and can determine directions for future AI Policy. Whether the incoming Trump administration will carry forward their recommendations remains to be seen.

Despite the infrastructure and thought leadership created since the executive order, the 2024 Republican Party platform vowed to "repeal Joe Biden's dangerous Executive Order that hinders AI Innovation, and imposes Radical Leftwing ideas on the development of this technology" (*The Republican Party Platform, 2024*, 2024). Upon taking office, President Trump repealed Executive Order 14110 (The White House, 2025c). At the time of writing, it is unclear which elements of the order might be incorporated into future policy. Republicans consider that the guardrails for AI focusing on equity and fairness hinder innovation and prioritize free speech and rapid AI adoption, potentially leading to a less regulated approach. This shift could result in the removal of anti-bias provisions and a reduced emphasis on mitigating AI risks (Alms, 2024). Shortly after revoking E.O. 14110, the Trump administration issued a brief Executive Order titled "Removing Barriers to American Leadership in Artificial Intelligence" with calls for the development in 180 days of an action plan to implement the policy established in Section 2 of the document: "It is the policy of the United States to sustain and enhance America's global AI dominance in order to promote human flourishing, economic competitiveness, and national security." (The White House, 2025d).

As discussed before under "The case for congressional action," congressional involvement, including bipartisan legislation, will also influence the final shape of federal AI governance.

U.S. Senate and House AI task forces

Executive Orders from the last three administrations, the failed legislative experience of the 118th Congress, Senate and House task forces, and changes in the judicial environment have set the stage for policy to be enacted. This work introduces a policy-agnostic way to do it.

In parallel to the executive branch activities, the 118th Congress released two major AI task force reports in 2024. In May, the Senate bipartisan AI Working Group released its "Driving U.S. Innovation in Artificial Intelligence," a self-titled roadmap for AI Policy outlining a strategy for advancing AI governance and innovation in the country (Hendrix et al., 2024). Led by Senate Majority Leader Chuck Schumer (D-NY), along with Senators Mike Rounds (R-SD), Todd Young (R-IN), and Martin Heinrich (D-NM), the group developed the roadmap following a series of educational briefings and nine AI Insight Forums held during the fall of 2023 including over 150 experts from industry, academia, and civil society. Organized into eight priority areas, it recommends allocating at least $32 billion annually to support non-defense AI innovation. This funding is intended to drive progress in AI research and development, enhance workforce capabilities, and establish ethical guidelines to ensure the responsible use of AI technologies. By focusing on innovation, investment, and ethical oversight, the roadmap aims to position the United States as a global leader in AI while addressing the technology's societal and economic implications. The report outlines legislative recommendations and areas for further exploration to advance AI governance and innovation in the United States. The Senate's approach to AI regulation in 2024 has drawn criticism for its lack of concrete action and failure to produce any draft legislation. Instead of proposing specific regulatory measures, the report offered a process for discussion without committing to policy outcomes or regulatory solutions (Green-Lowe, 2024). The measured approach, however, matches the pattern of how legislation tends to be built in Congress. One of the policy advisors we interviewed states: *"Almost no bill passes Congress the first time it's introduced. Typically, it requires two to five years for an idea to reach maturity in Congress and for every stakeholder to feel like they get a say. A lot of time, policy has been driven by the administration."*

In December 2024, the House Bipartisan Task Force on Artificial Intelligence released a comprehensive report that marked a significant step in shaping the nation's AI policy. Co-chaired by Representatives Jay Obernolte (R-CA) and Ted Lieu (D-CA), the task force was established in February 2024 by Speaker Mike Johnson and Democratic Leader Hakeem Jeffries. Its efforts resulted in a 273-page report containing 66 key findings and 89 recommendations, spanning 15 policy areas such as government use of AI, intellectual property rights, and energy consumption. The report underscored America's commitment to maintaining global leadership in responsible AI innovation while ensuring the implementation of necessary safeguards. It also highlighted AI's transformative potential in sectors like healthcare, where it could enhance diagnostic accuracy, streamline operations, and accelerate drug development. By addressing both opportunities and challenges, the task force aimed to set a balanced framework for harnessing AI's benefits

responsibly and sustainably (House Committee on Science, Space, and Technology, 2024). The reaction to the House report was still mixed, highlighting the tension between promoting AI innovation and addressing risks and harms (Iyer & Hendrix, 2024). While several praised the report's comprehensive approach, some critics – like with the Senate report - called for more concrete policy recommendations and stronger safeguards against AI-related risks. The report, however, was seen as a decisive step forward with an actionable vision for AI governance in the U.S., outlining what to regulate, who should regulate it, and how to do so effectively. It advocates for a sectoral, incremental approach to AI regulation and recommends further analysis relative to Federal preemption of State AI laws.

Both reports guide future congressional action on AI policy, emphasizing the importance of maintaining U.S. leadership in AI innovation while addressing potential risks and challenges. Common elements include promoting responsible AI development, recognizing the need for safeguards, and addressing national security, workforce development, and ethical considerations. Both reports resulted from extensive consultations with experts across various fields and covered a wide range of policy areas. The reports differ in their scope and approach. The House report is more comprehensive and emphasizes a sectoral regulatory approach, whereas the Senate report provides a broader policy roadmap. Along these lines, the House report goes deeper into specific sectors like healthcare and agriculture, while the Senate report takes a more general stance on AI governance. The Senate report set the stage for ongoing discussions, while the House report built upon these discussions and provided more detailed recommendations for future legislation. The two reports can be viewed as part of the same legislative development trajectory, laying the groundwork for action by the 119th Congress.

Let us take a closer look at what is happening in Congress.

Congressional perspectives

What legislative action is telling us

The 117th and 118th Congresses have shown significant activity regarding artificial intelligence legislation. During the 117th Congress (January 2021 to January 2023), lawmakers introduced 75 bills centered on AI and machine learning or that contained provisions specifically addressing these technologies. Six of these navigated the legislative process and were enacted into law, marking a step forward in addressing AI through federal policy (Harris, 2023).

This activity doubled in the 118th Congress (January 2023 to January 2025), with at least 150 AI-related bills proposed. (Brennan Center for Justice, 2024) However, none of these bills were enacted into law. One of our conversations with the Senate described an environment with *"a lot of fact gathering, and, frankly, a lot of fact gathering on facts that are already been gathered by other institutions."* In parallel and during the same period, 636 bills were introduced at the state level, of which more than one hundred were approved and are now State Law (Multistate.ai, n.d.).

In December 2024, Congress passed the FY 2025 National Defense Authorization Act (NDAA), which included the first AI provisions enacted by the 118th Congress (*Text - H.R.5009 - 118th Congress (2023-2024)*, 2024). However, the law is not regulatory and focuses instead on workforce development, creating AI education courses, and a Chief Digital Engineering Recruitment Officer within the Department of Defense. Within this scope, the act establishes pilot programs for AI in biotechnology, and manufacturing, and initiatives to improve AI usability. It also addresses AI safety and security, directing the identification of high-risk AI models and creating an AI Security Center. The NDAA emphasizes strategic planning, budgeting for AI data, and international collaboration through a multilateral AI working group (Sokler et al., 2024). The AI provisions from the NDAA primarily affect the Department of Defense (DOD), focusing on accelerating technology transitions from research and development to operational use. These provisions may also influence other national security and technology departments, such as the Department of Energy (DOE) and the Department of Homeland Security (DHS), through collaborative efforts in AI research, development, and implementation.

Additionally, it is noteworthy to look at the following five pieces of proposed AI legislation from the 117th Congress (MacCarthy, 2023). They were launched in 2021, following a recommendation from the lengthy "Investigation of Competition in Digital Markets" report from the House Judiciary Committee (U.S. Congress. Committee on the Judiciary Subcommittee on Antitrust, 2020). Starting in 2019, the committee performed a bipartisan top-to-bottom review of the market, focusing on the state of competition online and examining Amazon, Apple, Facebook, and Google business practices and their effect on the U.S. economy and democracy. It then assessed the ability of existing antitrust laws and competition policies to address market power and anti-competitive conduct in digital markets. While none were made into law, the report and associated legislative action significantly impacted the landscape of antitrust law regarding large tech companies. It highlighted potential anti-competitive practices and provided a framework for future legal action, new regulations, and increased scrutiny of mergers and business conduct in the digital market. The Federal Trade Commission (FTC) and the Department of Justice (DOJ) have been actively pursuing antitrust cases against major tech companies, including ongoing investigations and lawsuits against Amazon, Apple, Google, and Meta (parent company of Facebook) based on the practices identified in the report (American Economic Liberties Project, n.d.).

- Ending Platform Monopolies Act - **U.S. Congress**: *H.R. 3825, 117th Cong. (2021).* The Ending Platform Monopolies Act sought to eliminate conflicts of interest arising from dominant online platforms' concurrent ownership or control of an online platform and certain other businesses. It faced strong pushback from major tech firms, who argued it would stifle innovation, hamper consumer choice, and disrupt popular online services. Critics also worried about unintended consequences for small businesses reliant on large platforms' scale and reach and the viability of enforcement. The narrow committee passage signaled bipartisan skepticism about its sweeping restrictions, and it ultimately failed to advance, expiring with the adjournment of the 117th Congress.

- ACCESS Act - **U.S. Congress**: *H.R. 3849, 117th Cong. (2021).*

The ACCESS Act sought to mandate data portability and interoperability among major tech platforms. Critics raised privacy and security concerns, questioning how user data would be handled once transferred between services. Heavy lobbying from tech companies also stymied support, while debates arose about content moderation and editorial freedoms. It never reached a House floor vote and ultimately failed before the 117th Congress adjourned.

- American Innovation and Choice Online Act - **U.S. Congress**: *S. 2992, 117th Cong. (2021).*
  Despite strong bipartisan support, the American Innovation and Choice Online Act was derailed by concerns over user experience, product integration, and content moderation restrictions. The bill prevented discriminatory practices by large platforms, prohibiting them from giving preference to their products or services (self-dealing). Coupled with formidable tech industry lobbying, lawmakers worried about the bill's disruptive potential. Divisions within the parties over its scope and impact further stalled progress. Ultimately, these factors prevented the measure from reaching a Senate floor vote, and it expired with the adjournment of the 117th Congress, ending its prospects.

- Platform Competition and Opportunity Act - **U.S. Congress**: *H.R. 3826, 117th Cong. (2021).*
  Critics argued that the bill's stringent restrictions on large platforms' acquisitions threatened innovation and undermined startup opportunities, as it could deter venture capital and prevent smaller firms from being acquired. Tech companies opposed the measure, citing harm to user access to free or low-cost services. Despite bipartisan support and a narrow committee vote, the legislation ultimately lacked the momentum for a full House vote and died with the close of the 117th Congress.

- Open App Markets Act - **U.S. Congress**: *S. 2710, 117th Cong. (2021).*
  The Open App Markets Act promoted competition by reducing the gatekeeper power of dominant app store operators. It met resistance due to concerns about consumer security, privacy risks, and disruptions to the digital economy. Critics argued that permitting third-party payment systems and sideloading might inflate costs, compromise security, and hamper innovation. Despite broad bipartisan backing and a decisive committee vote, the legislation stalled without a Senate floor vote.

The failure of these five bills highlights deep-seated industry influence and persistent disagreements over how to balance competition with consumer protection.

Quantitative analysis of AI-related bills introduced in the 118th legislature

We used the Brenner Center for Justice Artificial Intelligence Legislation Tracker as our source for the federal AI bill inventory and the Congress.gov APIs to access detailed bill texts and associated metadata (Brennan Center for Justice, 2024; Library of Congress, n.d.). Applying

Machine Learning techniques, we identified and clustered patterns from the thousands of pages covered by these federal bills.

What follows is a quantitative analysis of the 150 AI-related bills introduced by the 118th Congress, revealing several key insights. The number is significantly higher than in previous sessions, indicating growing interest and concern among lawmakers regarding AI governance. Two-thirds of these bills have been introduced by Democrats (99 of the 150), suggesting a greater focus on AI policy within the Democratic Party.

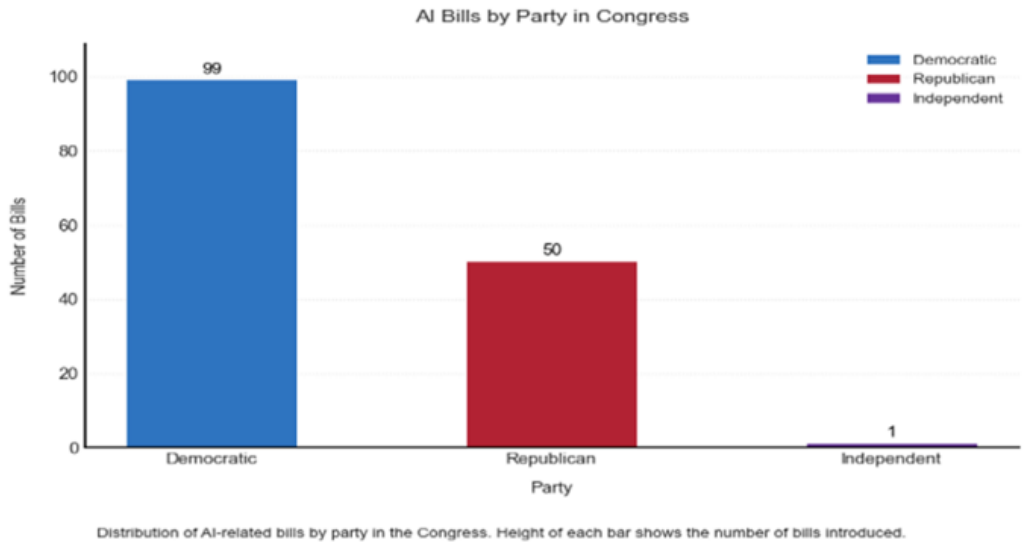

Figure 8 Bill sponsorship by party in the 118th Congress

Sponsorship clusters in the Democratic (72) and Bipartisan (54) categories indicate significant cross-party support for several initiatives despite the partisan divide on many issues. The total number of sponsors varies widely across bills, indicating diverse levels of support and coalition-building for different AI initiatives. Some bills have many sponsors but a relatively balanced partisan ratio, suggesting specific AI issues attract support from both parties.

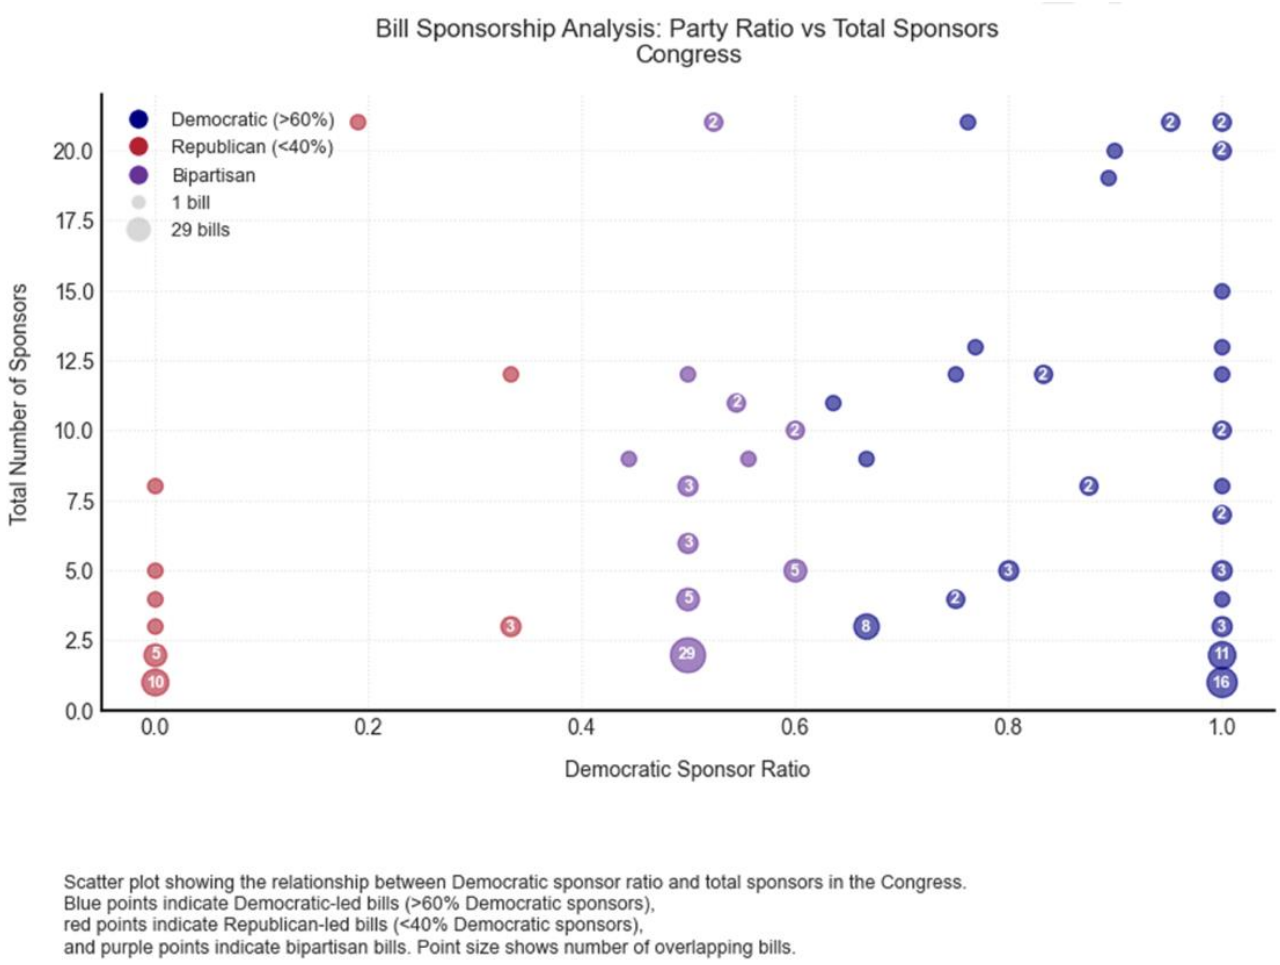

Figure 9 Bipartisanship analysis for AI-related bills in the 118th Congress

Members from 36 states introduced AI bills, indicating that this is not a localized issue. During the 2024 legislative session, at least 45 states, Puerto Rico, the Virgin Islands, and Washington, D.C., introduced AI bills at the state level. Of these, 31 states, Puerto Rico, and the Virgin Islands adopted resolutions or enacted legislation. (National Conference of State Legislatures, 2024)

Senators Mike Rounds [R-SD], co-chair of the Senate AI Caucus and member of the Senate's bipartisan AI working group, and Edward Markey [D-MA] sponsored the most AI bills in the Senate, with seven each. Representative Ted Lieu [D-CA], co-chair of the House's Bipartisan Taskforce on Artificial Intelligence, sponsored five bills, followed by representatives Clay Higgins [R-LA], Sheila Jackson Lee [D-TX], and Anna Eshoo [D-CA] with three sponsored bills each.

The Energy and Commerce Committee has been assigned the most AI bills in the House (20), followed by the Science, Space, and Technology Committee (15). The Commerce, Science, and Transportation Committee has been assigned the most AI bills in the Senate (33), followed by the

Homeland Security and Governmental Affairs Committee (8), and the Health, Education, Labor and Pensions Committee (8).

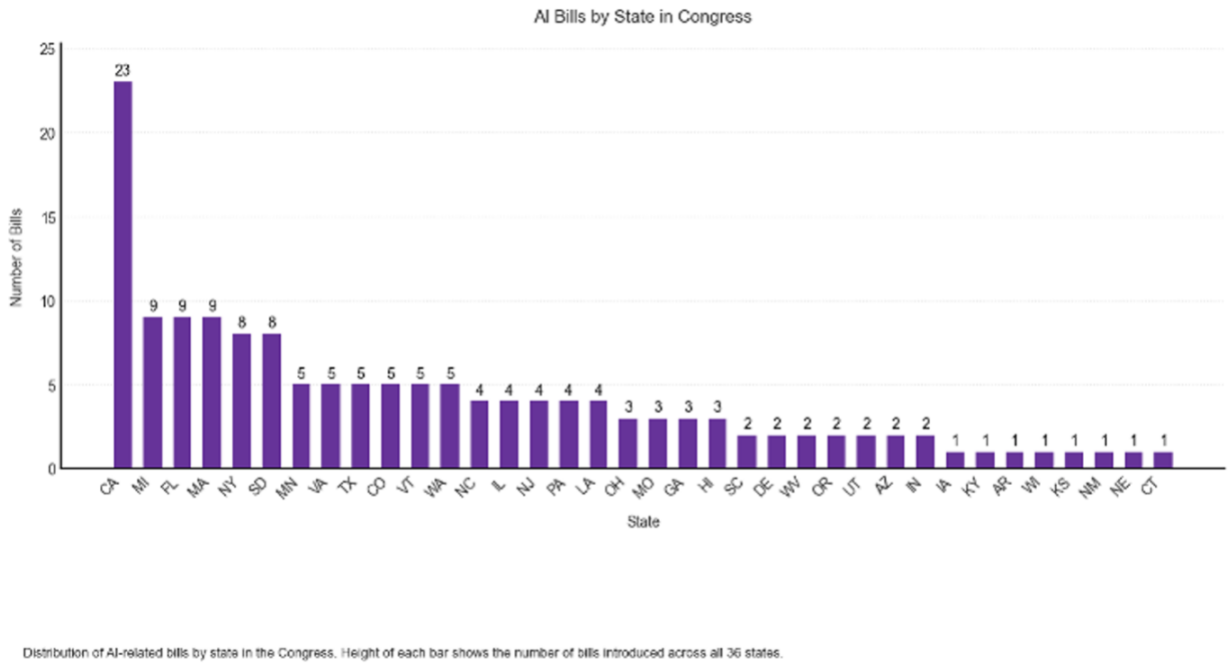

Figure 10 AI bill sponsorship by state during the 118th Congress

Overall, the metadata analysis indicates a growing focus on AI governance, with participation by Congress members from most states. The Democratic Party is taking the lead but with significant bipartisan support. Even though the increase in AI-related bills reflects members' newfound interest in the topic, our analysis showed that a few members continue to set the tone for the policymaking process.

The textual analysis of the 150 bills reveals the policy focus of each. The analysis of the "Primary Attributes in Congressional Bills" highlights the top emphasis lawmakers assigned to the bills. **Responsible and Ethical AI** emerges as the most common primary policy attribute, with 71 of the 150 bills mentioning it. This underscores a growing legislative concern for the ethical development and application of artificial intelligence. **National Security** follows with 24 bills, reflecting the importance to lawmakers of defense and security issues. In contrast, topics that

The analysis reveals a Congress focused on AI ethics, national security, and public interest, among other issues. This focus aligns with Democratic Party priorities. However, the frequency of policy attributes does not indicate bipartisan consensus. The in-depth interviews will surface possible paths forward.

dominate public discourse, such as **Antitrust** and **Copyright**, appear as primary attributes in only 5 and 4 bills, respectively.

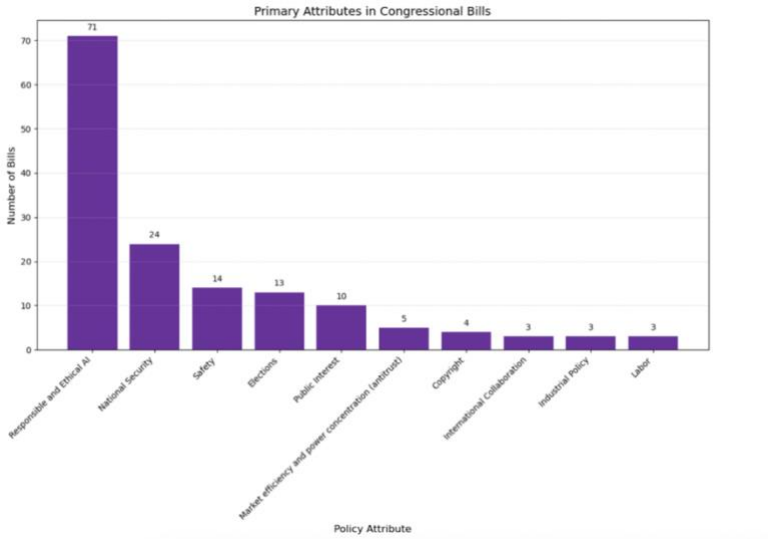

Figure 11 Primary policy attributes (150 AI bills from 118th Congress)

The “Top 3 Policy Attributes in Congressional Bills” analysis provides an expanded view by considering three main policy attributes from each bill. This broader perspective brings attention to issues like **Industrial Policy** and **Public Interest**, which gain prominence when secondary attributes are factored in. **Privacy** and **Data Protection** remain low-priority areas, highlighted in only two bills. This aligns with the U.S. legislative trend of deferring to state-level leadership on privacy issues rather than enacting federal regulations.

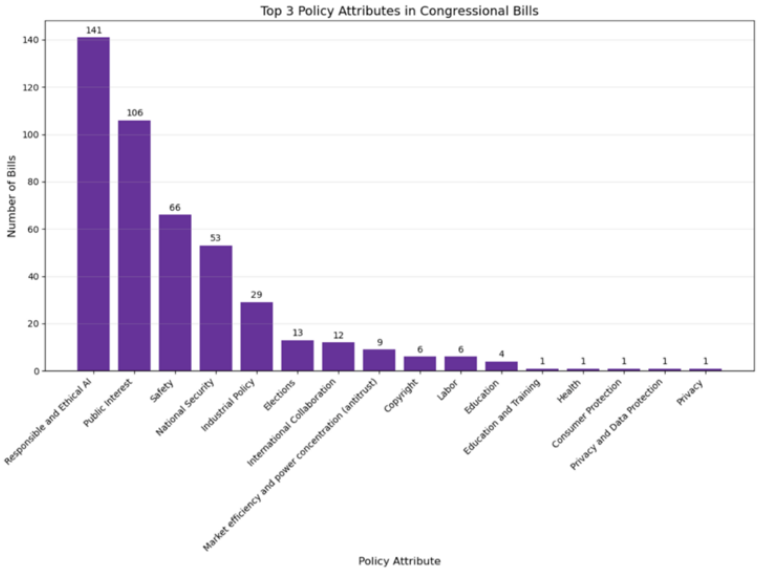

Figure 12 Top 3 Policy attributes (150 AI bills from 118th Congress)

When examining both scenarios, the legislative priorities of the 118th Congress become clear. Together, these analyses reveal a Congress grappling with AI ethics, national security, public interest, and a growing array of other critical issues. Lawmakers placed significant emphasis on the ethical development and use of AI and policies that serve the public interest. This focus aligns with Democratic Party priorities, which is unsurprising given that Democrats authored two-thirds of the bills.

However, the frequency of a policy attribute's appearance does not necessarily indicate bipartisan consensus. Qualitative interviews provide deeper insight into areas of alignment, as discussed later in this text. For example, key terms like "safety" have become less favored in Republican circles, where "guardrails" is the preferred terminology. National security and global leadership - particularly concerning the U.S. energy matrix - emerged as pivotal topics during the late Biden administration and the early days of the new Trump administration. These areas are expected to remain central to legislative efforts in the 119th Congress.

Congressional voices

Despite the volume of activity, little progress was made. No legislation was enacted, and only 17 out of the 150 bills were reported out of committee and entered the legislative calendar. To better understand the state of AI policy, we conducted extensive interviews with congressional staffers and members of Congress. These conversations revealed five key findings illuminating the complexities, priorities, and challenges shaping AI governance in the United States.

For many, regulation is less about fostering innovation and more about ensuring safeguards. The discourse often frames regulation as a tradeoff between innovation and safety rather than a synergistic tool to achieve both goals.

Our proposed model aims at achieving both.

One of the most contentious debates in AI policy revolves around whether regulation hampers innovation. Opinions are divided, often reflecting broader political alignments. Republican lawmakers view regulation skeptically, emphasizing concerns about stifling innovation and economic growth. On the other hand, Democrats, while more supportive of certain regulatory measures, approach the issue with varied motivations. For many, regulation is less about fostering innovation and more about ensuring safeguards to mitigate risks. Regardless of party affiliation, the discourse often frames regulation as a tradeoff between innovation and safety rather than a synergistic tool to achieve both goals, albeit to varying extents.

A staffer in the House of Representatives tells us: *"Pretty different opinions in Congress now on AI innovation and safeguards. It really needs to be both. AI will be embedded in everything. So, we need to pursue innovation, but also innovation in nontraditional ways, defense innovation, and public interest use cases, transport, energy, infrastructure, etc."* A House staff member adds: *"We want to do this gently, in a way that we're not allowing things to get out of control, but at the same time, we're not hampering innovation on AI."* A career staff member concludes: *"An approach to AI […] does not make regulation a goal in and of itself, right? And thinking through carefully what that actually means […] the first objective of every single U.S. policymaker should be to make sure that the U.S. has the most advanced, sophisticated, and rapidly developing AI ecosystem."*

A second notable theme that emerged is the perception in Congress that the absence of comprehensive congressional action on AI policy benefits business interests, particularly those of technology companies. This inertia creates an environment where businesses can operate with minimal constraints, leveraging the lack of regulation to pursue rapid development and deployment of AI technologies. While this benefits industry players in the short term, it raises concerns about long-term societal impacts and the role of government oversight. This finding was inconsistent with the view shared in the industry interviews, as we explored later in the text.

We heard from a former Senate staffer: *"It is in the interest of companies to do nothing. I mean, you could argue that they want the innovation stuff, but I think the status quo is perfectly amenable to what these companies would want."* A House staffer reflects that *"the most powerful people in the world aren't [in] Congress. They're the leaders of Deep Mind at Google; they're Sam Altman with ChatGPT."* The Senate staffer concludes that *"there's been some capture in this in the Senate when it comes to AI policy, just given the outstanding influence, or outsized, rather, influence, of AI companies and big tech firms."*

A lack of urgency for Federal action was a third theme. Speaking with a staffer serving on one of the Congressional Committees, we heard that *"it is a bit of like, okay, well, let's wait to see what happens there [in Europe] because folks are lacking […] a vision or imagination for what a regulatory landscape could look like, and what impact that has on the different types of stakeholders"* and a policy advisor reminded us that *"Congress is a democratic institution where every Member has the right to vote, [this] challenges the technocratic way of thinking about the technology. The most challenging part is coming up with policies that are appealing to a wide array of members. That may not be the best policy, but it's the only way to make progress."*

Despite AI's rapid advancement, policymakers feel little urgency to implement federal regulations. This delay is not due to a lack of understanding or innovative ideas; rather, it reflects deeper issues. The concept of AI regulation is still perceived as immature, with no consensus on what effective governance should look like. Political will remains insufficient, and the fragmentation of viewpoints contributes to a cautious approach. For many policymakers, the absence of immediate crises linked to AI reinforces the belief that regulation can wait. A staffer described the issue: *"Where Congress is most comfortable, is […] reacting to […] a rupture or big problem, where, when we're projecting forward […] it's not great, and even where we have a record but the issue is too complicated, or there's not […] enough consensus, it's hard to act, especially when we're*

talking about vote margins of just a couple of points or votes. […] We're waiting until there's something sort of discrete to you to fix, and mostly in a watching position."

The next strong theme reflected how industry engagement dominates policymaking. A congressional staffer warned: *"Bills do not usually pass the Congress over the wishes of major actors – that could be the administration, major industry groups, they could make it painful for you to move anything,"* and another interview added clarity stating that *"most of the AI policy, quite frankly, has been consolidated or is deeply controlled by leadership in both houses. […] There's a huge amount of lobbying and industry and money behind AI policy."*

Close engagement with the tech industry is a hallmark of congressional efforts on AI policy. Across the board, legislators acknowledge the need for industry involvement, whether as leaders or key supporters of any regulatory initiatives. Industry representatives often present their role as informative, offering insights into AI's technical complexities and potential impacts. However, this collaboration is not without criticism. The outsized influence of Big Tech on policymaking raises concerns about bias and the prioritization of corporate interests over public welfare. While some offices remain skeptical, most agree that any future regulatory framework will require significant industry input.

A staffer in the House ponders, *"I find it hard to believe that if something were to actually come to the floor, it wouldn't have already gotten buy-in or feedback from the actual industry. […] The leaders in AI in this country […] I mean, Apple, Google, Microsoft […] they're all leading the charge anyway."* At the same time, a Senate policy advisor tells us: *"Industry support or anti-support is crucial, maybe the most important honestly."*

The last of the five themes from our congressional interviews was the challenges to legislative progress faced by the 118th Congress. Despite growing discussions about AI governance, the 118th Congress has not passed any AI-related bills. A House Staffer reflects that *"members were almost afraid to try to push one solution when we know that any one solution could have actual consequences."* Our discussions revealed that several factors contributed to this legislative gridlock:

- **A Generally Unproductive Congress:** When asked why no AI-related bills passed in the 118th legislature, many interviewees initially pointed to the overall lack of productivity of the 118th Congress, which enacted only 3% of introduced legislation into law, compared to 7% during the 117th and 116th sessions (GovTrack.us, 2025). As one staffer stated, "*Congress right now is the least productive that it has ever been in its entire history*." However, digging deeper reveals a series of other factors.
- **Geopolitical Considerations**: Concerns about maintaining U.S. supremacy in AI and competing with China often overshadow domestic regulatory discussions. The concerns of a House Staffer are emblematic of this issue: *"I think there is always kind of a commitment there to the extent that […] we're not allowing things to get out of control, but at the same time, we're not hampering innovation on AI and progression on AI in this country and look back in three years and realize, wow, China, one of our biggest and most*

powerful adversaries, is now here with AI, and they're well-surpassed anything we've done on."

- **Political Polarization**: Deep divisions within Congress and the consolidation of power among leadership hinder consensus-building. *"I think anything can be polarized at this point,"* a frustrated staffer vents as she admits that *"it depends on what [leadership] wants."*
- **Focus on Nonregulatory Initiatives**: Many policymakers prioritize initiatives like research funding and voluntary guidelines over binding regulations. A congressional staffer confesses that *"[The] National AI Initiative Act [passed in 2020] is nonregulatory, focused on research standards and workforce development. Since that time […] we didn't do much. […] Now, there are dozens and dozens of bills from dozens of members."*
- **Fragmentation of Efforts**: With numerous competing initiatives and a lack of clear leadership, the path to actionable policy remains unclear.

The findings underscore the complexity of advancing AI policy in the United States. While regulation is viewed through a polarized lens, industry influence, geopolitical considerations, and legislative dynamics further complicate progress. Moving forward, fostering consensus, clarifying leadership, and balancing innovation with safeguards will be essential to building a robust and future-ready AI governance framework.

Contrary to the often-repeated notion that "nobody in Congress understands the technology," congressional members' expertise and access to experts in the field have improved over the last couple of years. Through sustained efforts in education, bipartisan dialogues, and task forces, lawmakers are becoming better equipped to handle the complexities of AI policy. One of our interviewees provided examples: *"We have a lot of great resources […] we can pull over people from the administration all the time to give us briefings or from industry groups […] we have the Congressional Research Service, the Government Accountability Office, one thing that my office wants to reinstate is the Office of Technology Assessment here on the hill. […] We have the congressional AI caucus […] and congressional ex-Staff Association […] to educate staffers and members on AI policies, AI issues that are coming up, and new technologies […] The task force was a really great resource that we had some really great briefings there, brought in some great subject matter experts."*

Where Congress and Industry could come together

According to our interviews, Congress and industry stakeholders share areas of alignment on key AI policy priorities, even though they often remain unaware of these shared perspectives. To different extents, both groups recognize the importance of establishing a regulatory framework that ensures safety, accountability, and trust without stifling innovation. Industry leaders, for example, frequently emphasized the value of clear federal standards to avoid the inefficiencies of fragmented state regulation, a concern echoed by congressional staffers who acknowledged the challenges of navigating a patchwork system. There is a solid alignment in that both the industry

and members of Congress prioritize and see the value of preserving, protecting, and extending the U.S.'s global leadership in AI.

Another point of agreement is the need for public-private collaboration. Industry representatives highlighted the importance of government-provided targets and clear regulatory benchmarks, reflecting their willingness to engage proactively in crafting solutions. Similarly, congressional staff underscored the value of advisory bodies and task forces that bridge technical expertise and legislative goals. Both sides support public-private partnerships to foster innovation while addressing critical societal risks like bias, safety, and transparency.

Targeted research and development (R&D) can drive innovation to address socio-technical challenges. Policy and regulation alone will not close this gap; actual technical invention is needed. Historically, the U.S. government has invested in research via its National Laboratories, grants to the thriving American university system, and partnerships with research bodies in the private sector. Those investments were a topic of consensus across Congress and the industry during our interviews. By investing in R&D that explicitly addresses social, ethical, and safety challenges, the government and the industry can unleash innovation that mitigates risks and creates a foundation for sustainable progress while maintaining a competitive edge.

Area of Agreement	Congress Perspective	Industry Perspective
Driving American Leadership in AI	Congress has focused on non-regulatory actions. To bolster U.S. leadership, policymakers and industry should collaborate on science, R&D, education, chips, quantum, and infrastructure. These efforts build on current momentum and foster trust, laying the groundwork for future regulatory measures.	The industry advocates for investment and workforce development, including expanding visa programs to enhance its talent pool. It also urges the government to prioritize infrastructure improvements and strengthen supply chains, recognizing their vital role in sustaining operations and maintaining competitiveness.
Focus on use cases and sectoral capacity	Congress emphasizes the need to update sector-specific legislation, enhance technical expertise, and clarify agency oversight for AI use cases, ensuring appropriate regulation across sectors like finance.	The industry supports targeted sectoral regulation, reliance on established expertise, and focus on practical test cases. It expects regulatory agencies to evolve alongside technology rather than create new frameworks for each innovation.
Little appetite for a new agency	Congress shows little interest in creating a new federal agency for AI oversight. While many support centralizing efforts—particularly for general-purpose AI—they differ on which existing department should take the lead. Several point to the National AI Initiative Act from 2020 and acknowledge that the AI Safety Institute has performed well within its mandate, though few advocate for significantly expanding its authority.	The industry is wary of creating a new regulatory agency but acknowledges the need for a broad, integrated approach to AI governance. Growing respect for NIST's work on standards reflects this perspective. There is alignment around the need to coordinate efforts in support of R&D, workforce development, infrastructure, and national security while retaining existing agencies for sector-specific oversight and enforcement.
Need to address AI energy requirements	There is bipartisan potential for addressing AI-related energy demand, which is expected to grow across all energy types. Efforts could involve legislation or regulatory action via the Department of Energy and the Federal Energy Regulatory Commission.	Big Tech, concerned about hitting limits in model development and scaling, prioritizes addressing these challenges. There are opportunities for partnership with local authorities in areas like water management, data center construction, and real estate development.

Figure 13 Examples of areas of alignment between Congress and the industry ecosystem (from interviews)

However, there remains a disconnect in communication and priorities. Industry voices often advocate for flexibility and innovation-focused policies, whereas congressional actors stress the need for robust safeguards to uphold public trust. Sometimes, this divergence is more about

emphasis than substance, suggesting an opportunity for collaboration. By building structured platforms for dialogue, Congress and industry can move beyond perceived differences and craft policies that balance innovation with accountability. An industry observer reflects on areas of consensus and bipartisan agreement: the *"main one is China, […] the U.S. should be the AI leader in the world, and that the U.S. cannot allow China to catch up or take over. In terms of how that's implemented, there's different viewpoints on what's most effective. So, some people are like 'Oh, maybe we should restrict some open source because that's what China is using for their most advanced models.' But then others are saying, 'Well, actually, you don't want any restrictions because they limit the ability of U.S. companies to run fast' […] policy areas that are bipartisan include things like making access to energy and data center build-out in the U.S., more readily available and expediting that."*

Area of Inconsistent Views	Congress Perspective	Industry Perspective
Privacy	Congress sees little chance of federal privacy laws advancing, with Democrats more open to it than Republicans.	Views privacy as critical but prefers alignment with existing GDPR principles to maintain global consistency. Resistant to disruptive new laws.
Intellectual property and copyright	Divided on IP and copyright, with representatives who advocate for content creators supporting action, while others preferring to wait for court outcomes.	The industry advocates for fair use and taking a wait-and-see attitude given the cases currently in court.
Market Competition and Power Concentration	Supports maintaining a competitive marketplace and fostering innovation, though opinions vary on action against Big Tech.	Big Tech advocates for government support to sustain global leadership, whereas smaller players favor competition, criticizing Big Tech for market dominance.
Workforce and skills	Workforce and skills development should be straightforward topics given the existing focus on education and R&D. However, debates such as those over H1B visas, highlight its political complexity.	The industry narrative focuses on augmenting human capabilities through co-pilot systems - rather than replacing employees - revealing sensitivity around workforce displacement.
Role of the States	The Congress acknowledges that the states are stepping in due to federal inaction, rather than a desire to lead AI regulation.	State-level efforts are a source of fragmentation and complexity. It doubts states' capacity to manage AI's intricacies and believes federal action can provide the uniformity and consistency needed for effective regulation and support of AI development.
Deep fake pornography and Child Sexual Abuse Material	Congress is likely to support measures targeting harmful AI use cases, such as deepfake pornography and child sexual abuse material, due to their focus on children's safety.	The industry broadly supports restrictions on high-risk AI applications and favors clear safety frameworks to mitigate risks. However, there are concerns about direct responsibility and liability.

Figure 14 Examples of areas where more alignment between Congress and the industry ecosystem is still needed (from interviews)

Technology policy requires an understanding of both: A bilingual fluency in science and mastery of the humanities. An industry insider reflected on this struggle: *"People in AI companies don't tend to know about policy. I'm […] passionate about policy and AI. […] In my team at [redacted] I was the only person from a policy background. Everyone had technical backgrounds, and they were brilliant at what they did, […] but what we do is hard, too. You can't just know computer science and intuit policy like it's something that you can pick up overnight. […] A lot of people who do AI don't know much about the world around, and they're fantastically good at coding, and so when they're thinking about risks […] they're not coming to it with a particularly nuanced view.*

[…] People who speak policy don't speak tech. […] People who speak tech don't really speak policy." Implementing effective policy requires bridging this chasm.

The industry is optimized to push the boundaries of technology and achieve market success. Democratic governments are designed to uphold public trust and protect societal interests. This contrast need not be a source of conflict. Instead, it offers a unique opportunity for collaboration.

The technological prowess of the private sector far outpaces that of the public sector, particularly in fields where rapid innovation and significant resources drive progress. This asymmetry reflects a fundamental reality: While industry is optimized to push the boundaries of technology and achieve market success, democratic governments are designed to uphold public trust and protect societal interests. These divergent mandates often create tension, as businesses prioritize profitability and shareholder value while governments emphasize safety, equity, and ethical considerations. However, this contrast need not be a source of conflict. Instead, it offers a unique opportunity for collaboration. Recognizing these complementary characteristics, a well-structured governance model can harness the strengths of both sectors.

Part III – A New Governance Model

After setting the context and introducing conceptual frameworks in Part I, we highlighted the results of our qualitative and quantitative research in Part II. Part III of the paper is dedicated to our proposal for a new Dynamic Governance Model.

This part describes the three components of the model:

- Public-private partnerships for the creation of evaluation standards.
- A market-based solution for audit and compliance.
- A system of accountability and liabilities set by legislatures, existing executive agencies, and the courts.

We have included an abbreviated example of implementing the model to drive one specific policy goal. This section concludes with a word of caution for those considering implementing the ideas in this manuscript and technology policy in general.

A path forward

The rapid adoption of artificial intelligence and its role as a transformative general-purpose technology has brought the United States to a pivotal moment. While the country has historically embraced innovation with a light regulatory touch, AI advancement's unique pace and scale demand a recalibration of this approach. The challenge is clear: How can we maintain the dynamism of technological progress while ensuring AI's safe, ethical, and equitable deployment? Addressing this challenge requires a new governance model that balances innovation with regulation, fosters collaboration between public and private sectors, and builds public trust in the potential of AI. This section outlines a strategic path forward, focusing on practical mechanisms to align technological development with societal needs and preserve America's leadership in this critical domain.

The American system, while imperfect, has consistently demonstrated its ability to drive technological and economic growth. This system, rooted in a combination of entrepreneurial spirit, public and private investment, and an open market, has fostered innovation across industries.

The American system, while imperfect, has consistently demonstrated its ability to drive technological and economic growth. This system, rooted in a combination of entrepreneurial spirit, public and private investment, and an open market, has fostered innovation across industries. From the early days of the Industrial Revolution to the rise of the Internet and modern digital technologies, the U.S. has maintained its leadership by balancing economic incentives with strategic policy interventions. Innovation has been a crucial driver of American economic growth and contributed as much as half of the nation's economic growth in the 20th century. This technological progress has led to new industries, products, and services that have fundamentally

transformed the American economy and improved living standards (The White House, 1995). The information technology (IT) sector has been a powerful driver of economic growth in recent decades, with an outsized contribution to the U.S. economy, creating high-paying jobs, driving exports, and spurring innovation across various sectors (Atkinson, 2022).

Like previous general-purpose technologies such as electricity and the Internet, AI could reshape industries, create new opportunities, and enhance the quality of life. A growing consensus among policymakers, industry leaders, and academics is that safeguards are essential to ensure AI applications' safe and ethical deployment. These safeguards aim to build public trust and minimize harm, yet the specifics of their implementation remain a matter of debate, as we have seen in the previous sections.

Technological innovation and legislation operate on fundamentally different timelines. Innovation often occurs rapidly and unpredictably, driven by breakthroughs, market dynamics, and entrepreneurial experimentation. By contrast, the legislative process is slow and intentionally designed to ensure thorough deliberation, consensus-building, and stability. This inherent disparity creates challenges in regulating fast-moving technologies like AI, where static regulatory frameworks can struggle to keep pace with the speed of development and deployment. The rapid evolution of AI demands a dynamic approach to governance that allows for adaptability and continuous refinement while preserving the measured oversight required to safeguard public interests.

A strictly ex-post approach to regulating AI, relying solely on antitrust actions and litigation, is insufficient to address this technology's unique challenges and catastrophic risks. While antitrust enforcement and legal recourse play a vital role in correcting abuses and ensuring accountability, they are inherently reactive. Such mechanisms often come into effect only after harm has occurred, making them ill-suited to prevent high-stakes risks like safety failures, systemic biases, or widespread societal disruption. Moreover, the complexity and pace of AI innovation can outstrip the ability of courts and regulators to respond effectively. At the same time, ex-ante regulation, while proactive, has its limitations. Static regulatory frameworks can struggle to keep up with rapid and unpredictable technological innovation, risking either obsolescence or overreach that stifles progress.

These phenomena underscore the need for a dynamic regulatory approach that combines the foresight of ex-ante measures with the adaptability to evolve alongside technological advancements. Such a system would complement existing ex-post mechanisms, creating a comprehensive framework to promote innovation and accountability.

Continued technological innovation is important to maintain the United States' economic and geopolitical leadership and address the pressing challenges inherent in current AI models. From the energy-intensive demands of large-scale computation to issues of alignment, bias, and safety, technological advancements are critical to ensuring AI systems are both sustainable and ethical. These challenges highlight the need for targeted investments in research and development, fostering breakthroughs in energy-efficient algorithms, bias mitigation techniques, and enhanced

model transparency. A well-crafted tech policy and regulatory framework must act as both an enabler and a safeguard, incentivizing the development of technologies that address existing limitations while ensuring accountability and public trust. By embracing a dynamic and collaborative approach to regulation, incorporating public-private partnerships and ongoing dialogue, policymakers can create an ecosystem where safety and innovation coexist.

The proposed dynamic governance approach brings government and industry together while combining the foresight of ex-ante measures with the adaptability needed to address technological advancements. Coupled with existing ex-post mechanisms, this creates a comprehensive framework to promote innovation and accountability.

Evolution in the governance of AI is needed, leveraging the strengths of both industry and government. Industry brings technical expertise, innovation capacity, and real-time insights into technological trends, while the government safeguards public interests through oversight, accountability, and policy direction. By integrating these complementary roles, such a scheme can adapt to the rapid pace of AI advancements, allowing regulatory frameworks to evolve in tandem with technological developments. This collaborative approach could take several forms, including joint standard-setting initiatives, co-development of best practices, and establishing regulatory sandboxes to pilot new technologies under controlled conditions. It also requires continual feedback and iteration mechanisms, such as advisory councils or working groups that bring together stakeholders from diverse sectors. Furthermore, embedding transparency and public engagement in this process will be crucial for building trust and ensuring that the framework remains responsive to societal needs.

A new model for collaboration and accountability

To address the challenges and opportunities presented by artificial intelligence, we propose a Dynamic Governance Model centered on structured collaboration between government and industry. At the heart of this model is creating an entity tasked with setting clear standards in partnership with the private sector. These public-private partnerships would go beyond mere consultation, fostering active collaboration in developing, implementing, and iterating on policies that align innovation with societal priorities. A key feature of this model is the incorporation of commitments by industry players underpinned by a robust accountability scheme, including clear delineation of liability and mechanisms for independent audits.

The extra-regulatory model was designed to be modular, light, and of general application. The modular nature of the model allows it to be implemented in phases, starting from different entry points. It is light enough to be built on top of and complement existing legislation and judicial sectoral frameworks, reducing the implementation burden on Congress and the industry ecosystem. Its general applicability allows policymakers to decide what use cases to prioritize (a risk-based approach, sectoral phased implementation, focus on larger companies, and gatekeepers are examples of policymaking and political decisions that can be made at implementation time).

This proposal requires decisive congressional action, including a review and potential expansion of statutory authorities to support the framework. By enacting clear policies, Congress can enable adaptive regulation that steers AI innovation in a safe direction and safeguards the nation's leadership in this critical enabling technology. This balanced approach ensures that while the transformative potential of AI is fully realized, its deployment remains aligned with ethical standards and public trust.

Figure 15 Dynamic Governance Model

The NTIA's efforts, featured in its March 2024 Artificial Intelligence Accountability Policy Report highlighted in "The Executive Order legacy" section, underscore the federal government's appetite to foster an AI accountability ecosystem. The NTIA has laid a foundation for ensuring trustworthy AI systems by emphasizing independent evaluations, standard-setting, and transparent disclosures. These initiatives resonate with our proposal for a dynamic and collaborative governance model, as both approaches highlight the importance of public-private partnerships, adaptive regulation, and the integration of technical standards to address AI's societal impact. Both frameworks advocate for accountability inputs—such as audits, disclosures, and documentation—as pillars of responsible AI governance. Likewise, we agree with the need for tiered risk-based approaches and the involvement of multiple stakeholders across the AI lifecycle. Our model diverges in its emphasis on extending public-private partnerships to include commitments enforced through a novel accountability scheme incorporating liability delineation

and audit mechanisms. This extension creates a dynamic, self-reinforcing system that evolves alongside technological advancements.

The Dynamic Governance Model

The Dynamic Governance Model begins with an essential role for the government: setting clear policy goals. As outlined in the “AI Policy Considerations” section, these goals should balance innovation with the public interest, ensuring that AI systems contribute positively to economic growth while safeguarding societal values. With these goals established, the model is made of three interconnected steps. First, Evaluation Standards are set via public-private partnerships that reflect technological feasibility, policy goals, and societal priorities. These standards create a benchmark for responsible AI development and deployment. Second, a market ecosystem for audits and compliance must be created, wherein independent entities assess adherence to the established standards. The process of audits and compliance ensures that AI systems are safe and ethical and meet transparency and reliability requirements. Finally, a robust system of accountability and liability must be codified and enforced by legislatures, regulatory agencies, and courts. Together, these components form a closed feedback loop where the insights gained from implementation, compliance monitoring, and enforcement inform the continuous improvement of the evaluation standards. With active participation from the government, industry, and civil society, we embed trust into this cycle. The dynamic model ensures adaptive governance that evolves alongside technological advancements. We will now look at the details for each step.

Evaluation standards

The success of the Dynamic Governance Model hinges on establishing robust and adaptive evaluation standards. These standards must be developed through collaborative public-private partnerships that draw on the expertise of industry leaders, academic researchers, policymakers, and civil society organizations. Existing models, such as the U.S. Artificial Intelligence Safety Institute (AISI) under NIST, provide a framework for AI safety and risk management. However, as the NTIA Accountability Report points out, these efforts often lack enforcement power and standardized methodologies for independent evaluations (NTIA, 2024a). Public-private partnerships should set actionable, measurable, and enforceable standards to address this gap.

International models, such as the EU’s regulatory sandboxes and the OECD’s Observatory of Public Sector Innovation (OPSI), demonstrate how experimental environments can foster regulatory innovation. While these initiatives promote valuable collaboration, they also highlight the need for scalability and broader adoption. The implementation of these novel models is still in its early stages. The regulatory sandbox work in Europe predates the EU AI Act, which now provides these environments with a more substantial legal basis. Spain is pioneering nationwide regulatory sandboxes and represents an important case for those looking for empirical evidence of success, but the jury is still out (Future of Life Institute, n.d.). The NTIA report stresses the

importance of integrating accountability mechanisms—such as mandatory disclosures and formal evaluations—into these models to ensure they are more than voluntary exercises. The EU AI Act embeds a similar approach under the concept of “regulatory learning.”

Importantly, evaluation standards should encompass technical and non-technical dimensions of AI accountability and reflect technological feasibility, ethical considerations, and societal values. The NTIA’s emphasis on information flow, including standardized documentation and disclosures, supports this goal. By requiring that developers and deployers provide clear, consistent information on AI system design, capabilities, risks, and limitations, stakeholders across the AI ecosystem can better assess compliance and trustworthiness.

The implementation of evaluation standards requires a flexible, iterative approach that can accommodate technological evolution. While this work does not aim to provide a fully developed implementation plan, we recommend several key steps. First, as suggested by the NTIA, a structured public-private dialogue platform could be formalized to ensure ongoing collaboration in updating standards. This platform would allow the sharing of best practices, emerging risks, and lessons learned from real-world applications. One attempt at creating such a platform is the “Promoting United States Leadership in Standards Act of 2024,” introduced in December 2024 – **U.S. Congress**: *H.R. 10281, 118th Cong. (2024)*. A bipartisan effort by the Senate, also in December 2024, introduced “A bill to establish the Artificial Intelligence Safety Review Office in the Department of Commerce, and for other purposes” – **U.S. Congress**: *S. 5616, 118th Cong. (2024)* (Nihill, 2024). These examples demonstrate Congress’s efforts to codify initiatives originally established through executive orders. As discussed before, the Supreme Court decisions that limit agency discretion make incorporating these initiatives into statutory language critically important.

Second, pilot programs, like those pioneered by the EU, should be launched in sectors where AI adoption presents significant risks and opportunities, such as healthcare, finance, and critical infrastructure. These pilots would provide crucial feedback on the practical implementation of Evaluation Standards and inform their refinement.

Finally, the NTIA report highlights the importance of developing and certifying independent evaluators and auditors to foster consistency and trust. While audits and compliance will be discussed in the next section, evaluation standards must be directly tied to these downstream accountability mechanisms. Only through regular, independent evaluation can these standards remain relevant and effective in mitigating risks while fostering innovation.

A valid criticism of a public-private partnership for setting evaluation standards is the risk of regulatory capture by powerful industry actors. Several strategies can be employed to mitigate this risk. First, ensuring balanced stakeholder representation by including diverse groups such as civil society, academia, consumer advocates, and smaller businesses reduces the risk of dominance by a few large players. Second, establishing independent oversight mechanisms—such as third-party audit bodies—enhances impartiality in enforcement. The following section will cover the oversight mechanisms in more detail. Transparency is critical, with public consultations

and mandatory disclosure of interactions between regulators and industry representatives fostering trust. Robust conflict-of-interest policies, including cooling-off periods for regulators transitioning to industry, help limit undue influence. Legislative oversight, judicial review, and use of decentralized standards-setting bodies, such as NIST or IEEE (Institute of Electrical and Electronics Engineers), can further prevent regulatory capture by distributing decision-making authority. Adopting performance-based regulations, which focus on outcomes rather than prescriptive rules, limits industry influence over specific technical requirements. Finally, periodic reviews and sunset clauses ensure that policies remain relevant and effective over time.

Strategy	Mechanism	Expected Benefit
Balanced Stakeholder Representation	Inclusion of civil society, academia, SMEs, and consumer advocates	Reduces risk of dominance by large tech companies, ensuring diverse perspectives
Independent Oversight	Creation of independent third-party audit and compliance bodies	Ensures impartial enforcement and reduces industry influence on compliance
Transparency	Public consultations, open meetings, and mandatory disclosure of interactions	Deters lobbying and builds public trust in the process
Conflict of Interest Policies	Cooling-off periods for regulators; strict rules on industry ties	Prevents revolving door effect and limits undue influence
Legislative and Judicial Safeguards	Regular oversight by a congressional committee; judicial review of actions	Ensures checks and balances and provides recourse in case of capture
Decentralized Standards Setting	Involvement of multi-stakeholder bodies like IEEE or NIST	Distributes decision-making power, reducing concentrated influence
Performance-Based Regulations	Focus on measurable outcomes rather than prescriptive requirements	Encourages innovation and limits rules favoring specific incumbents
Periodic Review and Sunset Clauses	Mandatory periodic reviews and inclusion of sunset clauses	Ensures policies are updated and remain effective over time

Figure 16 Anti-capture strategies for public-private partnerships

A market ecosystem for audits and compliance

Creating a robust market ecosystem for audits and compliance will represent a key improvement over current voluntary commitments, which, while valuable, often lack enforcement mechanisms and consistency (The White House, 2023a). One can draw lessons from other well-established industry-led compliance frameworks to build an effective ecosystem.

A relevant example is UL (Underwriters Laboratories), one of the largest and oldest independent safety science companies (Kagan, 2023). For over a century, UL has set industry benchmarks for

product safety, certifying products ranging from consumer electronics to medical devices. Its model demonstrates how an independent, industry-recognized certification body can enhance safety, build trust, and facilitate market access. Similarly, CE marking in Europe (Conformité Européenne) ensures that products meet health, safety, and environmental protection requirements, enabling free trade across the European Economic Area (European Union, n.d.). UL and CE illustrate how structured, industry-wide compliance mechanisms can contribute to market trust and safety, enabling innovation and competition. Both entities have faced criticisms regarding conflicts of interest, lax standards, and, in the case of UL, monopolistic practices. CE's self-declaration approach raises concerns about rigor and enforcement. Recently, UL has made a series of organizational and process changes in preparation for its 2024 IPO. CE has updated emerging technologies and sustainability regulations while extending recognition in the UK post-Brexit. Both systems continue to evolve.

There are also significant parallels to be drawn from accounting standards. Systems like U.S. GAAP (Generally Accepted Accounting Principles) and IFRS (International Financial Reporting Standards) have established a common framework for financial reporting across industries and countries. The Securities and Exchange Commission (SEC) acts as one of the regulatory bodies overseeing the standards and mandates transparency through extensive disclosures and regular audits. This approach ensures that financial information is reliable and comparable, fostering trust in financial markets. Significantly, accounting standards are periodically revised to reflect changes in business environments, demonstrating the need for iterative updates to standards as technology evolves. GAAP and IFRS face ongoing criticisms regarding their complexity, potential for manipulation, and ability to reflect a company's financial position accurately. However, it would be hard to imagine a world without accounting standards and the associated economic chaos, reduced transparency, impact on business operations, and impeded global financial cooperation and development.

In much the same way, a market ecosystem for AI audits and compliance requires standardized evaluation protocols, clear regulatory oversight, and mechanisms to ensure that standards evolve. The standardization, compliance, transparency, and continuous improvement principles underpin product certification and financial reporting frameworks and offer a blueprint for creating a similar ecosystem in the AI space.

Building from the NTIA's recommendations, the following steps can be considered for establishing this ecosystem:

1. **Standardized Audit Criteria**: Like evaluation standards, public-private partnerships should define standardized criteria for AI audits, specifying what aspects of an AI system (e.g., safety, fairness, privacy, transparency) must be evaluated and how. These criteria should be publicly accessible and regularly updated to reflect technological advancements and emerging risks.
2. **Certification and Oversight Bodies**: Independent certification bodies, like UL or financial accounting firms, should be established to conduct audits. These bodies would be responsible for certifying AI systems against the established standards. Sectoral

regulatory agencies could oversee these bodies, ensuring accountability and adherence to auditing protocols.

3. **Auditor Qualification and Certification**: To ensure high-quality audits, auditors must be certified based on clear qualifications and expertise in the technical and ethical aspects of AI. A national auditor training and certification program could be launched in collaboration with universities and industry partners.
4. **Mandated Disclosures and Reporting**: As with financial reporting, companies deploying high-risk AI systems should be required to disclose key information about their systems' design, testing, and deployment. Standardized disclosure formats, such as model cards and datasheets, would enhance transparency and enable stakeholders to assess compliance more effectively.
5. **Sector-Specific Pilots**: Pilot programs could be launched in high-risk sectors (e.g., healthcare, finance, autonomous vehicles) to test the feasibility of audit protocols and identify best practices for implementation. Insights from these pilots would inform the broader rollout of audit requirements across industries.
6. **Iterative Updates**: Just as accounting standards are regularly revised, AI audit criteria should be updated based on feedback from auditors, regulators, and industry stakeholders. This iterative process will ensure that standards remain relevant and robust in the face of rapid technological change.
7. **Federal Support for the Ecosystem**: Federal agencies should provide programmatic support to accelerate the development of a vibrant market for AI audits. This support includes funding research on auditing methodologies, developing technical audit infrastructure, and facilitating public-private partnerships. Additionally, federal procurement rules could require government suppliers and contractors to comply with certified AI audit standards, further incentivizing industry participation.

By embedding these elements into the market ecosystem, the governance model ensures that AI systems are subject to continuous scrutiny and improvement, fostering trust among all stakeholders—developers, deployers, regulators, and the public—while encouraging innovation under a framework of accountability and transparency. This process is inherently long-term, requiring incremental progress over several years and likely maturing over decades. However, this should not deter us from taking the necessary first steps to initiate the development of this ecosystem. Early efforts can provide the foundation for a mature, sustainable ecosystem. These include establishing pilot programs, defining baseline criteria, and creating a framework for auditor certification.

Congress, especially the Senate, has demonstrated an interest in exploring the intersection of standards and how to enforce compliance. Here are a few examples of bills that have been introduced in the Senate during the last couple of years in this area:

- The "TEST AI Act of 2024"**, US Congress:** *S. 3162, 118th Cong. (2024)* mandates the creation of AI testbeds for trustworthy system development, involving interagency collaboration, risk assessments, and classified facilities to prevent AI misuse. It

emphasizes red- and blue-teaming methodologies for vulnerability evaluations and aims to strengthen national security and critical infrastructure protection.

- The “AI Research, Innovation, and Accountability Act of 2023”, **US Congress:** *S. 3312, 118th Cong. (2023)* promotes AI innovation while ensuring accountability through transparency, risk management frameworks, and generative AI system disclosures. It mandates transparency reports for high-impact AI systems, recommends best practices for government AI adoption, and encourages robust AI oversight to mitigate risks in sensitive sectors.
- The “VET AI Act”, **US Congress:** *S. 4769, 118th Cong. (2024)* directs NIST to develop voluntary guidelines for AI assurance, focusing on internal and external evaluations of AI systems. It aims to enhance trust through meaningful assurances, safeguard privacy, mitigate harms, and support the adoption of trustworthy AI solutions via evidence-based standards and technical validation.
- The “Future of Artificial Intelligence Innovation Act of 2024”, **US Congress:** *S. 4178, 118th Cong. (2024)*, establishes and codifies an AI Safety Institute, promotes innovation via testbeds and capacity-building initiatives, and supports international AI standards and research collaborations.

Each bill involves NIST in developing AI standards, metrics, and methodologies to ensure safety and accountability and emphasize collaboration between public and private sectors, academia, and international partners for AI research and development.

On a side note, the concept of regulatory markets introduced in the last decade in academic and AI research lab circles presents an innovative approach to governance by addressing the weaknesses of traditional government-led regulation and industry self-regulation (Clark & Hadfield, 2019). This model focuses on governments setting clear regulatory goals while private regulators compete to deliver effective solutions that meet these objectives. By fostering competition, the approach drives innovation in regulatory technologies, enabling oversight mechanisms to keep up with the rapid evolution of AI systems. Regulatory markets tackle this issue by delegating responsibilities to private entities, incentivizing them to create advanced tools and methodologies for auditing, monitoring, and safeguarding AI systems. These private regulators operate under governmental supervision to maintain accountability while allowing for a more flexible and adaptive regulatory framework. Despite its promise, the regulatory market model faces challenges, particularly in maintaining robust government oversight and preventing regulatory capture. Nevertheless, it offers a dynamic pathway to align AI development with societal safety and ethical norms (Carvao, 2024a).

Accountability and liability

The third and final step of the Dynamic Governance Model is the establishment of a robust system of accountability and liability. This step addresses the question of who is responsible when AI systems cause harm or fail to meet established standards and has historically proven to be the most challenging to implement. Accountability mechanisms are critical to ensuring that AI’s

promises (innovation, efficiency, and better decision-making) do not come at the expense of safety, fairness, or human rights.

A helpful reference for understanding the importance of liability frameworks is the experience with social media platforms under Section 230 of the Communications Decency Act. Section 230's liability shield enabled the rapid growth of the Internet and Web 2.0 by protecting platforms from being held liable for user-generated content. While this provision facilitated unprecedented innovation and economic growth, it also attracted criticism for allowing platforms to evade responsibility for harmful content. This example illustrates the delicate balance between encouraging innovation and ensuring accountability—a balance that must now be carefully navigated in the context of AI (Carvao, 2023).

Liability in the AI ecosystem introduces new complexities due to the distributed nature of AI development and deployment. Unlike traditional product liability, where responsibility can often be traced to a single manufacturer, AI systems are developed and deployed through a multi-layered supply chain. Key actors include **developers** (the entities that create and train AI models), **deployers** (those who integrate and implement AI models into products or services), and **end users** (Individuals or organizations that operate AI-driven systems in specific contexts). The NTIA Accountability Report emphasizes the need to ensure accountability across this entire value chain, highlighting that liability should not rest solely on one actor but be shared proportionally based on each actor's level of control and responsibility. For example, developers may be held accountable for transparency and design flaws, while deployers could be responsible for ensuring that AI systems are used appropriately and safely. The NTIA report identifies several areas where changes in law and policy are required to support AI accountability:

1. **Privacy**: Stronger protections are necessary to prevent AI systems from infringing on individuals' privacy. These include clear rules on data usage, retention, and sharing and requirements for transparent disclosures about AI system behavior.
2. **Copyright**: With AI models increasingly trained on vast datasets, including copyrighted materials, new legal frameworks are needed to clarify AI developers' and content creators' rights and obligations.
3. **Safety**: AI systems that pose risks to physical safety—such as autonomous vehicles and medical devices—should be subject to sector-specific safety regulations, including mandatory pre-deployment testing and certification.

Topics such as privacy and copyright reform are among the most challenging to address and, as previously noted, represent areas lacking a prevailing normative consensus. The model's modular, lightweight, and broadly applicable design becomes especially valuable when confronting these difficulties. Rather than beginning with these complex issues, one can first focus on areas of more explicit agreement to build trust in the model. Once trust is established, the same approach and process can be gradually extended to more protracted issues.

Effective enforcement of accountability and liability frameworks will require a coordinated effort across multiple institutions, each playing a distinct and complementary role. **Congress** will be

instrumental in establishing the legislative foundation for AI accountability, including codifying baseline liability standards and mandating sector-specific regulations where necessary. Additionally, Congress can create mechanisms for ongoing regulatory oversight to ensure these frameworks remain effective and adaptive as AI technology evolves. Even though no AI bills have been enacted into law, there has been active work during the 118th Congress in this direction. Here are several examples: the "AI Fraud Deterrence Act"**, U.S Congress:** *H.R. 10125, 118th Cong. (2024)* increases penalties for financial crimes such as mail, wire, and bank fraud committed using AI. The "Artificial Intelligence Civil Rights Act of 2024"**, US Congress:** *S. 5152, 118th Cong. (2024)* establishes protections against discrimination by computational algorithms in key areas like employment, housing, and healthcare and outlines transparency, individual rights, and enforcement mechanisms to safeguard civil rights against algorithmic bias and harmful outcomes. The "NO FAKES Act of 2024", **US Congress:** *H.R. 9551, 118th Cong. (2024)* protects intellectual property rights related to individuals' voice and visual likeness, addressing unauthorized digital replicas created using AI. It includes provisions for civil action against violations and safe harbor rules for online services that promptly remove infringing content upon notification. Each bill includes provisions for enforcement, penalties, or legal recourse to ensure compliance.

Sectoral regulatory agencies, such as the FDA for healthcare and the SEC for financial markets, will oversee the implementation of AI-specific regulations within their respective domains. These agencies will be tasked with certifying high-risk AI systems before deployment and ensuring compliance with established safety and ethical standards. They will play a crucial role in maintaining public trust in AI systems used in critical sectors by providing targeted oversight.

Courts will also serve as a vital component of the enforcement ecosystem. As disputes arise, courts interpret liability standards and resolve conflicts, setting legal precedents that shape how accountability is distributed across the AI supply chain. Over time, judicial rulings will help clarify ambiguous areas in legislation and regulation, contributing to the development of a more robust and well-defined framework for AI governance.

Existing legal frameworks can also address challenges associated with AI. For example, the recent legal advisories issued by California Attorney General Rob Bonta highlight how state laws—including consumer protection, civil rights, competition, and data privacy laws—can be applied to AI systems. These advisories emphasize that AI does not exist in a legal vacuum and that businesses and developers are subject to longstanding legal principles. For instance, under California's Unfair Competition Law, practices like deceptive advertising of AI capabilities or the misuse of AI for fraudulent purposes are actionable. Similarly, civil rights laws protect against bias and discrimination perpetuated by AI systems, while privacy laws ensure that AI developers handle sensitive data responsibly. These measures underscore the potential to adapt existing legal standards to novel technologies like AI, providing a pathway for governance that leverages familiar structures while addressing emerging risks. Such an approach enables policymakers to bridge the gap between innovation and accountability without necessitating entirely new regulatory frameworks for every technological advance (State of California Department of Justice, 2025). In parallel, California has enacted new AI-specific laws requiring transparency in training

data and AI-generated content, safeguarding personal likenesses, ensuring disclosures in campaign materials, penalizing exploitative uses like deepfake pornography, and mandating licensed oversight for healthcare AI tools. These laws complement existing frameworks to address the unique risks of AI technologies. The advisories mark a pioneering effort by a state to provide explicit guidance on applying existing laws within the evolving context of AI technologies.

While not an exact comparison, other domains draw on established frameworks in fields like product safety certification, financial accounting, and insurance regulation, where robust systems of accountability and liability have been successfully developed and implemented over the years. We discussed product safety certification examples such as UL and regulatory frameworks like the European CE marking. We also touched upon Financial Accounting and systems like U.S. GAAP and IFRS, offering lessons in transparency and compliance enforced by regulatory bodies such as the SEC. In insurance regulation in fields like healthcare and automotive industries, companies are often required to carry liability insurance to cover potential damages caused by their products or services. This model can be adapted to AI, where high-risk systems might similarly require mandatory insurance or contributions to compensation funds to safeguard against unforeseen harm.

By building on these established approaches, the AI ecosystem can adopt proven mechanisms—such as codified liability rules, independent oversight, and incentivization of best practices—to foster a culture of responsibility and trust. These lessons inform the following steps that can help operationalize a system of accountability and liability:

1. **Codification of Liability Standards**: Congress should pass legislation establishing clear liability standards for AI systems, including rules for determining fault across the AI supply chain.
2. **Sector-Specific Regulations**: Regulatory agencies should develop detailed, sector-specific guidelines for AI deployment. These guidelines should include pre-deployment testing, certification, and post-deployment monitoring requirements.
3. **Insurance and Compensation Funds**: Companies deploying high-risk AI systems could be required to carry liability insurance or contribute to compensation funds designed to cover potential damages.
4. **Auditor and Certifier Accreditation**: Independent auditors and certifiers should be accredited by recognized bodies to ensure consistent and reliable evaluations of AI systems. These auditors are key in identifying potential risks and ensuring compliance.
5. **Incentivizing Best Practices**: Policymakers could create incentives for companies that adopt best practices in AI accountability, such as tax benefits or reduced regulatory burdens for certified systems. These incentives will encourage proactive compliance.

While building a comprehensive framework for AI accountability and liability is a long-term process, early steps such as establishing baseline liability standards and launching sector-specific pilot programs can set the foundation for a mature and effective system. Over time, this framework can evolve to reflect new technological developments, emerging risks, and lessons learned from real-world deployments.

Ultimately, ensuring accountability and liability in AI will require significant legal and regulatory advancements. Much of this must be codified into laws and statutes, clearly defining liability across the AI supply chain. Courts will play a crucial role in interpreting and applying these laws, setting precedents that refine liability standards over time. While this is a complex and long-term process, establishing a solid legislative and judicial foundation is essential for fostering trust, promoting innovation, and safeguarding societal well-being in an AI-driven future.

Implementation context

The proposed Dynamic Governance Model aligns well with the U.S. approach to policy implementation, integrating elements of antitrust, sectoral regulation, and litigation. It employs antitrust principles to facilitate ex-post corrections when market concentration or anti-competitive behavior undermines the public interest. Implementing sectoral regulations provides ex-ante safeguards, establishing clear rules and standards before any harm occurs. Furthermore, the model incorporates aspects of tort law by emphasizing liability throughout the AI supply chain, where courts will play a vital role in interpreting these standards and establishing legal precedents.

The Dynamic Governance Model aligns well with the U.S. approach to policy implementation, combining elements of antitrust, sectoral regulation, and litigation.

The safeguards generated by the model will mitigate risks so that the industry can continue to accelerate and innovate.

This hybrid approach reflects the U.S.'s ability to balance innovation with regulation, ensuring that technological progress is not stifled but guided responsibly. By adopting this model, the U.S. can enhance its geopolitical standing in AI leadership, outpacing global competitors such as China by demonstrating that democratic governance can foster innovation and accountability. Unlike top-down regulatory regimes, this model emphasizes transparency, fairness, and stakeholder involvement, which can strengthen trust in U.S.-led AI initiatives on the international stage.

Codifying liability standards, creating audit ecosystems, and fostering public-private partnerships will enable the U.S. to maintain its AI geopolitical, technical, and economic leadership. Furthermore, by collaborating with international partners on harmonized standards, the U.S. can build global coalitions for responsible AI governance, ensuring interoperability, enhancing trust, and promoting shared ethical values. Most importantly, this model promotes trust—both domestically and internationally—by demonstrating a commitment to ethical AI development, ensuring safety and fairness while fostering continued innovation.

As we find ourselves in February 2025, at the start of a new administration and the 119th Congress, the conversation surrounding technology and AI policy in the United States is beginning to reach a more advanced stage. One of our recent congressional interviews

underscores this evolving maturity: *"The educational phase has gone well, the private sector's […] R&D of what we need from the policy side of things has gone well, and now, to […] cap it off, the executive branch and some of the biggest players in the private sector are […] publicly stating that AI is going to continue to be a massive initiative. […] The culmination of all three of those things puts us in a good position"* to have a productive 119th Congress.

Implementation example: Improving data center energy efficiency

To illustrate the practical application of the Dynamic Governance Model, consider the pressing challenge of addressing the energy demands of AI systems and their data centers. As highlighted in the interviews with congressional and industry stakeholders, there is an alignment between the two on the need to address the rising energy consumption of AI technologies. Additionally, there is growing bipartisan consensus on the topic – albeit with different approaches on how to do so – making it a good candidate to test-pilot the Dynamic Governance Model.

Policy Goals and Objectives

The hypothetical policy goal is to enhance the energy efficiency of data centers housing AI systems. This objective might involve Congress passing legislation to:

1. Establish benchmarks and metrics for energy efficiency through the Department of Energy (DOE).
2. Incentivize innovation in energy-efficient hardware, algorithms, and data center design.
3. Create testbeds for evaluating and validating energy efficiency improvements.

The Dynamic Governance Model provides a structured framework to implement such goals fostering collaboration and accountability across stakeholders.

Evaluation Standards

Continuing with this simulated scenario, the Department of Energy would spearhead the development of evaluation standards for data center energy efficiency, establishing a public-private dialogue bringing together:

- **Regulatory Agencies**: Including DOE, NIST, and relevant offices within the Departments of Commerce and Energy (including expertise from National Labs).
- **Industry Stakeholders**: Representing Big Tech, startups, and the investment community (tech, real estate, and others) in a diverse set of resources, expertise, and interests, including representatives of civil society like labor unions and affected communities.
- **Independent Oversight**: Ensuring transparency and accountability in the standard-setting process and preventing regulatory capture.

Depending on the scope of the energy efficiency measures being targeted, the exact office that is in the lead on the evaluation standards may vary. In this particular case, it could be the Federal Energy Regulatory Commission (FERC) or the Office of Energy Efficiency and Renewable Energy (EERE). Strong input from NIST would be a necessity, given the rapidly evolving nature of these technologies.

This dialogue platform would operate over a defined timeline, producing actionable, consensus-driven standards. This would be an iterative process that accounts for the evolving nature of these standards. Unlike standards for other sectors, standards and metrics for AI models will change quickly as the technology advances. The DGM's dynamic nature allows it to account for these changes and regularly update these evaluation standards.

Audit and Compliance Framework

The established standards would start as a voluntary program and, with the support of Congress through enacting legislation, evolve to be subject to formal audits. These audits could require:

- Disclosure of energy efficiency data by AI companies.
- Pilot programs to validate the feasibility and effectiveness of audit protocols.
- Regular updates to audit criteria, informed by feedback from auditors, regulators, and industry participants.

After the initial pilot phase, federal procurement rules could mandate compliance with these standards, incentivizing widespread adoption and fostering a culture of accountability.

The ENERGY STAR voluntary labeling program administered by the Environmental Protection Agency (EPA) in partnership with the DOE represents a similar initiative for products like appliances, electronics, lighting, heating, and cooling systems (U.S. Environmental Protection Agency, n.d.). In 2010, the Government Accountability Office (GAO) exposed vulnerabilities in the program's self-certification process. In response, the EPA mandated third-party certification starting in 2011, requiring testing by recognized labs and oversight by accreditation bodies.

The data center energy efficiency program can eventually evolve towards an auditable program with enforceable compliance. Data center energy consumption has a material economic impact and a lower volume of projects when compared to a more consumer-product-oriented program like ENERGY STAR.

Accountability and Liability Mechanisms

If audits reveal noncompliance with the established standards, liability frameworks would come into play. Remedies could include:

- Withdrawal of research funding.
- Discontinuation of tax incentives.
- Financial penalties or fees.
- Mandated corrective actions.
- Public disclosure of violations.

Guidance on these remedies must be outlined in legislative frameworks, ensuring clarity and consistency in enforcement. They can be invoked as part of pre-existing statutes or will need to be approved by Congress and made into law to avoid issues with agency discretion. In case of disagreement, the power to assign liability and assess damages will continue to fall to the courts. In addition to explicit standards compliance audits, precedent-setting judgments will clarify expectations of who is responsible for what when not clearly defined in the statutes.

The use of existing laws and statutes or the creation of new ones was a topic of discussion during our interviews. In the words of a legislative staffer, *"It's […] an 80/20 approach, 80% what […] we have in place already that we can apply and then accepting the fact that there are new things we have to come up with, and that's kind of that 20% of something new, and that's the approach we've tried to take. […] That way, we're not wasting our time coming up with new things. […] It's the fine line between how much [new] is necessary."*

Implementation of the model, on the other hand, creates an opportunity for transparency and the sharing of best practices in a virtuous cycle.

Sectoral and Legislative Synergy

The DOE and other sectoral agencies may need to update their policies for data center energy use. Congress could further mandate energy efficiency requirements for AI developers and data center operators, backed by specific penalties for non-compliance. Such measures would ensure alignment between legislative objectives and sectoral implementation. This would have to be reflected in the affected agencies' statutory mandates as defined by Congress, ensuring enforceability.

Looking ahead with the Dynamic Governance Model

The scenario outlined above is far from fictional. Similar policy objectives have been proposed during the 118th Congress as part of the House's Department of Energy Artificial Intelligence Act of 2024 - **U.S. Congress:** *H.R. 9671, 118th Cong. (2024)* – and the Senate's Department of Energy AI Act - **U.S. Congress:** *S. 4664, 118th Cong. (2024)*.

The House bill focused primarily on research funding and appropriation. It commissioned the creation of a strategic plan, to be reported back to Congress in no more than one year, outlining

short-term and long-term goals and resource needs to advance applications of AI for science, energy, and national security. It was not regulatory.

The Senate bill established programs to advance AI research, infrastructure, workforce, and security within the DOE mission scope. It acknowledged the department's leading role in AI for national security, science, and energy, given their 17 National Laboratories, which employ over 40,000 scientists, engineers, and researchers with decades of experience in advanced computer science research. It directs the creation in three years of shared infrastructure to support the development of frontier AI models. Like the House bill, it also allows for partnerships with private entities. The bill introduces a new Office of Critical and Emerging Technology within the Office of the Under Secretary for Science and Innovation to oversee these activities and authorizes the appropriation of the necessary funds from fiscal years 2024 through 2032.

While broader in scope, the Senate bill in this example did not focus as much on compliance and accountability, which are key attributes of our proposed model. It is not regulatory. By complementing the research, investment, and partnership focus from the bills with auditable and enforceable standards, one creates an innovation-inducing effect, much like what was discussed in the "Regulation-induced innovation" section. New economic incentives (or liabilities) and engineering constraints will lead to out-of-the-box thinking and the creation of new solutions and designs.

The Dynamic Governance Model's versatility can be extended to diverse policy objectives, including addressing privacy concerns, market competition, and intellectual property challenges. It is designed to be uniquely suited to the challenges and ideas identified during the primary research and interviews with both Congress and the industry. In contrast to the perceived burdens of traditional regulation, the Dynamic Governance Model's incremental, adaptive nature makes it an invaluable tool for navigating complex, evolving policy landscapes.

In contrast to the perceived burdens of traditional regulation, the Dynamic Governance Model's incremental, adaptive nature makes it an invaluable tool for navigating complex, evolving policy landscapes.

This underscores the Dynamic Governance Model's potential to translate high-level policy goals into actionable, collaborative frameworks. By fostering innovation, ensuring accountability, and balancing stakeholder interests, this approach can address the multifaceted challenges posed by AI and its associated technologies.

A cautionary tale

The influence of Big Tech on the development of artificial intelligence policies is both a boon and a burden to the broader technological ecosystem. As innovation champions, these companies are undeniably essential to the U.S. economy and play a pivotal role in maintaining the U.S.'s global AI leadership. However, their dominance has also led to market distortions, raising critical questions about Congress' role in sustaining these dynamics and the implications for competition policy (Baer, 2025).

Congressional interviewees openly acknowledged the outsized influence of major technology firms on AI policymaking. One respondent noted the *"subtle machinery of the Hill"* through which Big Tech controls the legislative agenda. This influence is evident in Majority Leader Schumer's AI Insight Forums, where *"the likes of Elon Musk, Mark Zuckerberg, Bill Gates, and Sam Altman all in one room together with a few civil society advocates, sure, but across […] nine sessions; I don't think that the non-industry actors had a […] seat at the table, quite literally."* The resulting 2024 report, entitled 'Driving U.S. Innovation in Artificial Intelligence: A Roadmap for Artificial Intelligence Policy in the United States Senate,' devoted significant attention to pro-innovation strategies, with scant mention of risks like privacy and liability (United States Congress. Senate, 2024).

A letter to the Federal Trade Commission from a group of Senators *"calling out some concerns about consolidation in the AI industry, […] there was this pattern […] of big tech firms really involving themselves in the development of the AI field in pernicious but subtle ways,"* as described by a Senate staffer exemplifies how large firms consolidate their dominance while skirting regulatory scrutiny. Big Tech circumvents traditional antitrust oversight by hiring key employees from smaller AI startups and licensing their technology, further entrenching its market power (Matt O'Brien & Sarah Parvini, 2024). These actions are not anomalies but part of a broader pattern of equity deals, acquisitions, and strategic partnerships to absorb emerging competitors before they become a threat.

Despite bipartisan interest in holding these companies accountable, Big Tech's entrenched influence often leaves Congress in a reactive posture. The inability to enact changes to the antitrust laws highlights a systemic prioritization of industrial policy and national security over competition policy, allowing Big Tech to shape discussions around AI to its advantage. Our proposed Dynamic Governance Model seeks to leverage Big Tech's capabilities and strengths. However, it is vital to recognize the risks of over-concentration.

Ultimately, our governance model aligns with this nuanced approach, advocating for a system where Big Tech's power is harnessed under the supervision of government and civil society. However, this deliberate choice comes with the responsibility to ensure that policies do not merely

entrench existing power dynamics but create an equitable and competitive environment for all stakeholders in the AI ecosystem.

Conclusion

Envision a world ten years from now where the rapid rise of AI has reshaped society as profoundly as the stormy weather once reshaped the fate of the tugboats Montrose and T.J. Hooper. Coordinated global efforts steer AI's transformative power toward solving humanity's most significant challenges. Transparent governance and ethical standards enable breakthroughs in combating climate change, eradicating diseases, and enhancing education. Workers find fulfillment as automation takes on repetitive tasks, fostering creativity and economic growth. Societies thrive as inequalities shrink, guided by policies prioritizing equity and sustainability.

However, like the tugboats caught unprepared for the storm, another future looms. AI deepens societal divides without cohesive regulation, empowering monopolies, expanding surveillance, and embedding biases within opaque systems. Strained resources and escalating conflicts lead to stagnation and mistrust. Innovation, once a beacon of hope, becomes a source of instability.

The tugboats' lesson is clear: ignoring the benefits and risks of new technology can lead to catastrophe, but thoughtful preparation can avert disaster. Today, we stand at a similar crossroads. Our choices will determine whether we harness AI as a force for shared progress or face a storm of division and missed possibilities.

This paper has explored the need for a balanced and adaptive framework to govern artificial intelligence, emphasizing the unique challenges posed by this rapidly evolving technology. Beginning with analyzing shared goals and perspectives between Congress and industry stakeholders, we highlighted the convergence of interests around the need for trust, transparency, and safety in AI.

We then introduced the **Dynamic Governance Model**, a three-step approach that builds on public-private partnerships to establish evaluation standards, develop a market ecosystem for audits and compliance, and ensure accountability and liability across the AI supply chain. Each step draws lessons from existing regulatory models and incorporates insights from interviews with policymakers and industry leaders. The model proposes a practical path forward, balancing innovation with regulation, fostering global competitiveness, and embedding trust into AI systems.

In conclusion, the Dynamic Governance Model offers a balanced path forward for AI policy, addressing safety, accountability, and innovation concerns while fostering public trust. By drawing on established regulatory frameworks - antitrust, sectoral regulation, and tort law - and adapting them to AI's unique challenges, the model provides a comprehensive approach that leverages public and private expertise. The three-step framework ensures that AI governance remains adaptive, effective, transparent, and aligned with societal values.

At this pivotal juncture, the U.S. stands at a crossroads. Our choices today will shape the domestic AI landscape and the global competitive environment for decades. By adopting this model, the U.S. can lead by example, demonstrating that it can balance innovation with regulation without stifling technological progress. This moment represents a chance for the U.S. to establish itself as the trusted global leader in AI governance, offering a democratic alternative to more authoritarian approaches and ensuring alignment with shared international values.

As Shakespeare reminds us, *"What is past is prologue."* The lessons learned from previous technological revolutions have prepared us for this moment. Madison's words in *Federalist No. 51* - *"If men were angels, no government would be necessary"* - underscore the need for governance when power and technology intersect. Now, at this critical crossroads, we can take the right turn and create a governance framework that balances innovation with accountability and ensures that AI's transformative potential serves humanity's broader interests. By doing so, the U.S. will safeguard its future and build a foundation for global cooperation, trust, and sustainable progress.

Acknowledgments

This paper would not have been possible without the support of many individuals. I am deeply grateful to my colleagues at the Mossavar-Rahmani Center for Business and Government, particularly those in the Technology, Innovation, and Regulation Program, for providing a platform to explore the intersections of AI, governance, and innovation. The foundational ideas for this work, rooted in systems thinking and a focus on social impact, emerged from the Advanced Leadership Initiative at Harvard University and were enriched through discussions with faculty and students from Harvard and MIT. This manuscript benefited from the expertise of research assistants from the Kennedy School and members of the Harvard Undergraduate Machine Intelligence Community, they are my co-authors.

When mentioning specific names, there is always a risk of inadvertent omissions, for which I take full responsibility. The individuals acknowledged here provided invaluable support through coaching, mentoring, manuscript feedback, and numerous other contributions: Nicholas Ashford, Yochai Benkler, Antonio Claudio Buffara, Alejandra Castillo, Justin Hendrix, Eugene Kimmelman, Chris Lewis, Mark MacCarthy, Martha Minow, Christina Montgomery, Nathan Sanders, Jonathan Sallet, Bruce Schneier, T.L. Taylor, Jim Waldo, Thomas Wheeler, and Jonathan Zittrain. My 2023 research assistants helped shape the early ideas on the interaction of technology and democracy: Iman Alshawk, Wren Oppermann, and Andrew Steen. This group of Kennedy School master's students participated in a reading group to review the final draft: Jack Conlin, Brian Moscioni, Giulia Neaher, Casey Reimche, Lucas Schmuck, Darryl Slabe, and Jake Steckler. In addition to my co-authors, the following members of HUMIC contributed valuable technical expertise: Omer Mujawar, Aarna Pal-Yadav, Lily Xu, and Joshua Zyzak.

This working paper would not have been possible without candid discussions with 49 industry and congressional leaders, who will remain anonymous under our commitment to confidentiality.

Special thanks to my academic advisor, John Haigh, Lecturer in Public Policy at the Harvard Kennedy School, and the Senior Fellows in the 2024-2025 cohort, especially Gregory Makoff for the multiple discussions about data methods and analysis.

Thanks to my family and friends for putting up with my obsession with this topic over the past three years - I know I have been sounding like a broken record.

The authors used Grammarly Pro to assist with grammar and style. All AI-generated content was reviewed and revised by the authors to ensure accuracy and alignment with the research findings. The authors take full responsibility for any errors or omissions.

Appendix I – Approaches and Methods

The impetus for this work was developed during the principal author's term as an Advanced Leadership Initiative Fellow at Harvard University in 2023. After a career in technology and business, the author became increasingly concerned with polarization and radicalization and intrigued by technology's role in furthering these issues. Initially imbued with a techno-deterministic view of the problem, the first thrust was to try and fight a technology problem with technology – i.e., content moderation challenges would be addressed with better use of artificial intelligence to filter content, misinformation issues would be tacked with better content provenance and attribution schemes, and so on. Upon taking a holistic and systemic view of the problem, it became apparent that the socio-technical pieces of the puzzle were as important as the technical ones. This realization led to a journey of visiting the incentives that fuel the technology industry today and a view that reinvention of its business models and regulation is needed to allow for a safe and ethical future of the banner technology of the 21st century: AI. The question then became how to do it without harming innovation and, instead, how to use policy to drive innovation in a direction that each society deems normatively correct. The nexus of the work then became what policy approaches should be taken by the United States of America, the home of the largest AI labs and the biggest of the Big Tech firms. This working paper was enabled by a 2024-2025 fellowship with the Technology, Innovation, and Regulation Program at the Mossavar-Rahmani Center for Business and Government of the Harvard Kennedy School.

With research assistance from Master of Public Policy candidates at the Kennedy School and in partnership with the Harvard Undergraduate Machine Intelligence Community (HUMIC)—all co-authors here—we set out to examine U.S. digital tech policy decisions and explore a path forward that bridges the tech industry, as drivers of innovation, with Congress, as representatives of the American society. Following a broad literature review, instead of taking a normative position, we engaged in a discovery journey to challenge our preconceived ideas about the topic and let a picture emerge from our primary research. In addition to the sources cited in this work, we have conducted primary research in two areas:

1. **23** long-form, qualitative interviews with members of the U.S. Congress or their staff
2. **23** long-form, qualitative interviews with technology industry leaders working in AI

These non-attributable interviews, with signed confidentiality and consent forms, averaged thirty minutes to an hour and generated hundreds of pages of transcripts, allowing the complexity of the subject to emerge. When quoted in the text, they are marked in italicized font and may have been edited for style and brevity while keeping the original meaning of the quote. These interviews aimed to map each stakeholder's interests, agendas, and perceptions to identify overlaps and explore collaborative solutions.

For the congressional interviews, participants were selected based on seniority and experience in AI policy. We spoke to **25**[7] individuals, members of Congress, current and former staff working in members' offices, as professional committee staff, or as fellows from outside organizations. Everyone interviewed had several years – and in some cases, decades – of experience working in the field of Artificial Intelligence, with about half having policy expertise and the other half having domain-specific technical expertise. Topics of conversation included AI experience, AI policy views, experience working on AI in Congress, assessment of institutional and legal capacities of U.S. institutions to respond to AI development, overall dynamics in Congress, engagement with industry and other stakeholders, engagement with international stakeholders, and the state of AI policy in the U.S. today.

For the industry interviews, participants were selected based on seniority and expertise in AI technology. This cohort included **24**[7] industry leaders, encompassing investors, engineers, startup founders, researchers, academics, board members, and executives from AI companies. Participants were purposefully chosen to represent a broad variety of organizations, ranging from Big Tech to industry-aligned groups, startups, consulting firms, research labs, and more. All participants are engaged within the U.S. AI ecosystem and will be directly impacted by future U.S. AI policy. Additionally, most have significant global engagement in their professional roles.

Each interviewee possessed multiple years—some with decades—of experience in Artificial Intelligence, primarily from the business perspective, though many also brought technical, policy, or ethics expertise. Topics explored during the interviews included their professional experience with AI, perspectives on AI policy, insights into the AI ecosystem, key concerns, and challenges related to navigating AI policy uncertainty. Other areas of discussion involved internal procedures and working processes for addressing the unique aspects of AI, assessments of institutional and legal capacities of U.S. entities to respond to AI developments, interactions with the U.S. Congress, identification of legislative frameworks that could be effective in the future, engagement with international stakeholders, and evaluation of the current state of AI policy in the U.S.

While our work was being done, the 118th Congress of the United States was ending. During their tenure, 150 AI-related bills have been proposed, and none have been enacted into law. We used the Brenner Center for Justice Artificial Intelligence Legislation Tracker as our source for AI bill inventory (https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker). In parallel, 636 State Bills were introduced, and dozens have been approved and are now State Law (https://www.multistate.ai/artificial-intelligence-ai-legislation). In partnership with HUMIC, we used Machine Learning and Artificial Intelligence techniques to identify and cluster patterns in the thousands of pages covered by the 150 Federal bills. Appendix II includes the details of this work. The analysis of the state legislation is not in the scope of this project.

[7] The difference between number of interviews and interviewees accounts for interviews which included more than one person.

We also analyzed “public materials”, such as podcasts, public domain interviews, and public lectures. Our insights were drawn from the primary sources listed above, the literature review, and the machine learning analytical work of the legislative activity during the 118th Congress.

Appendix II – Analysis of AI bills (118th Congress)

The primary objective of this analysis is to evaluate the breadth and depth of legislative activity related to AI within the 118th Congress, providing insights into emerging policy priorities, areas of bipartisan consensus, and potential gaps in regulatory approaches. By examining the content, partisanship, and progress of individual legislation, we aim to better understand the current scope of proposed AI legislation during this period (January 3, 2023, to January 3, 2025) and assess how lawmakers responded and prioritized challenges posed by AI.

This appendix offers a structured analysis of:

- The key policy areas addressed by AI-related bills.
- Patterns of sponsorship across partisan affiliations and geographic regions.
- The progression and legislative success rates of AI-related proposals.
- The extent of alignment between legislative efforts and industry perspectives on AI governance.

Data Methodology

Data Collection

We began by utilizing the Congress.gov API to retrieve all bills introduced into the House and Senate from the 118th Congress. An example output from the API call is displayed in **Error! Reference source not found.**[8]. From the information provided by the API, we store the bill number, origin chamber, bill title, bill text, summary, and list of updates. The congressional API also includes its own broader policy area categorization that can be found on the API's website (Congress.gov, n.d.).

From this dataset, fetched bills were filtered down to only the 150 bills in the 118th Congress that were manually identified by researchers at the Brennan Center as AI-related or AI-adjacent (Brennan Center for Justice, 2024). For each identified bill, we gathered information about sponsorship patterns, committee assignments, and legislative progression. For each sponsor and cosponsor, we gathered district information, party affiliation, the day they sponsored, and whether they were an original sponsor or signed on after the bill was first introduced. This allowed us to assess the level of bipartisan cooperation and the likelihood of legislative success based on historical trends.

[8] At the start of the 119th Congress, the routes for api.congress.gov were changed. API calls are now done using gpo.congress.gov.

Consumer Safety Technology Act

Attribute	Details
Bill Number	4814
Congress	118
Title	Consumer Safety Technology Act
Type	HR
Origin	House
Policy Area	Commerce
Introduced Date	2023-07-20
Latest Action	Received in the Senate and Read twice and referred to the Committee on Commerce, Science, and Transportation.
Sponsor	Rep. Soto, Darren [D-FL-9]
Cosponsors	4 (Reps. Burgess, Trahan, Guthrie, Castor)
Committee Count	2
Actions Count	20
Amendments Count	null
Related Bills Count	null
Subjects Count	18
Summaries Count	3
Text Versions Count	4
Titles Count	19
Update Date	2024-11-09T05:06:33Z

Figure 17 Example of returned API request data for H.R.4814 in tabular form

Data Cleaning and Preprocessing

Although the congressional API provides a detailed list of how bills progress in both chambers of Congress, the descriptions used to represent the status and progression of the bills contain inconsistencies. For example, the first markup of a bill can be entered in the API as “Mark-up”, “Markup”, or “Mark up”, while still representing the same action. We used regular expressions to harmonize different descriptions that had the same outcome to categorize bills into ten levels of progress, starting from the bill being introduced (Stage 0) to the bill being signed into law (Stage 10):

Stage 0 – Introduced
Stage 1 – Referred to Committee
Stage 2 – Referred to Subcommittee

Stage 3 – Reported Out of Committee
Stage 4 – Placed on Calendar
Stage 5 – Floor Consideration
Stage 6 – Passed One Chamber
Stage 7 – Received by Second Chamber
Stage 8 – Passed Second Chamber
Stage 9 – Presented to President
Stage 10 – Signed into Law

It is worth noting that the legislative process is not always conducted in an interactive, easy-to-follow fashion. Whenever the last action on the bill did not match a specific stage, we referred to the most recent stage that fit the categorization above.

Two bills from the dataset did not match any of the regular expressions:

- S.B.3696, the DEFIANCE Act of 2024, passed the Senate and was “Held at the desk.” in the House, placing it at stage 7.
- S.B.2103, the Intelligence Authorization Act for Fiscal Year 2024, had the last action “By Senator Warner from Select Committee on Intelligence filed written report. Report No. 118-59. Additional views filed.” The second-to-last action was the bill being placed on the Senate Legislative Calendar on 6/22/2023, placing it at stage 4.

```
keywords = {
   0: [r"\bIntroduced\b", r"Referred to (?!.*Subcommittee).*Committee"],
   1: [r"\bReferred to Subcommittee\b", r"\bSubcommittee\b"],
   2: [r"\bMark-up\b", r"\bMarkup\b", r"\bMark up\b"],
   3: [r"\bReported\b", r"ordered to be reported", r"\bYeas\b", r"\bNays\b"],
   4: [r"Placed on calendar", r"\bcalendar\b", r"Calendar No\."],
   5: [r"Discharged", r"Considered under suspension of the rules", r"Measure laid
before Senate"],
   6: [r"Passed Senate", r"Passed House", r"pass the bill"],
   7: [r"Received in the Senate", r"Received in the House"],
   8: [r"Senate agreed to", r"House agreed to"],
   9: [r"Presented to President"],
   10: [r"Signed by President"]
}
```

Figure 18 Code snippet of regular expressions used to preprocess progress data.

Handling Missing Data

All bills were searched for missing data. Some bills were identified with missing summary fields and were supplemented with manual reviews in addition to the analysis of the bill text. Additional bills had duplicate metadata in the sponsor and cosponsor list, which were pruned to ensure accuracy.

Tagging Legislation

Selection of Policy Attributes

To identify legislative trends, we categorized bills into major policy domains, including national security, privacy and consumer protection, workforce impact, AI governance frameworks, and liability/accountability measures outlined in the list below. These categories were developed based on legislative text analysis and consultation with existing taxonomies of AI policy concerns as, for example, those from Figure 5.

- Market efficiency and power concentration (antitrust)
- Safety
- Responsible and ethical AI
- National security
- Industrial policy
- Public Interest
- Labor
- Copyright
- International collaboration
- Elections

NLP-Based Bill Classification

To enhance the categorization of AI-related legislation, we implemented Natural Language Processing (NLP) techniques to classify bill text on a first pass-through before conducting manual verification. This approach allowed us to systematically identify key policy themes, dominant topics, and relevant stakeholders embedded in legislative texts. The NLP methods applied included Term Frequency-Inverse Document Frequency (TF-IDF), Latent Dirichlet Allocation (LDA), and Named Entity Recognition (NER).

TF-IDF was used to identify first the most important keywords and rank policy focus areas. We tokenized each bill's text into individual words and phrases and removed stop words (e.g., "the," "and," "is") to improve relevance. After removing stop words, we applied TF-IDF weighting to score terms based on their frequency in a specific bill relative to their frequency across all bills in the dataset. Higher TF-IDF scores indicated terms uniquely important to each bill, revealing key legislative concerns.

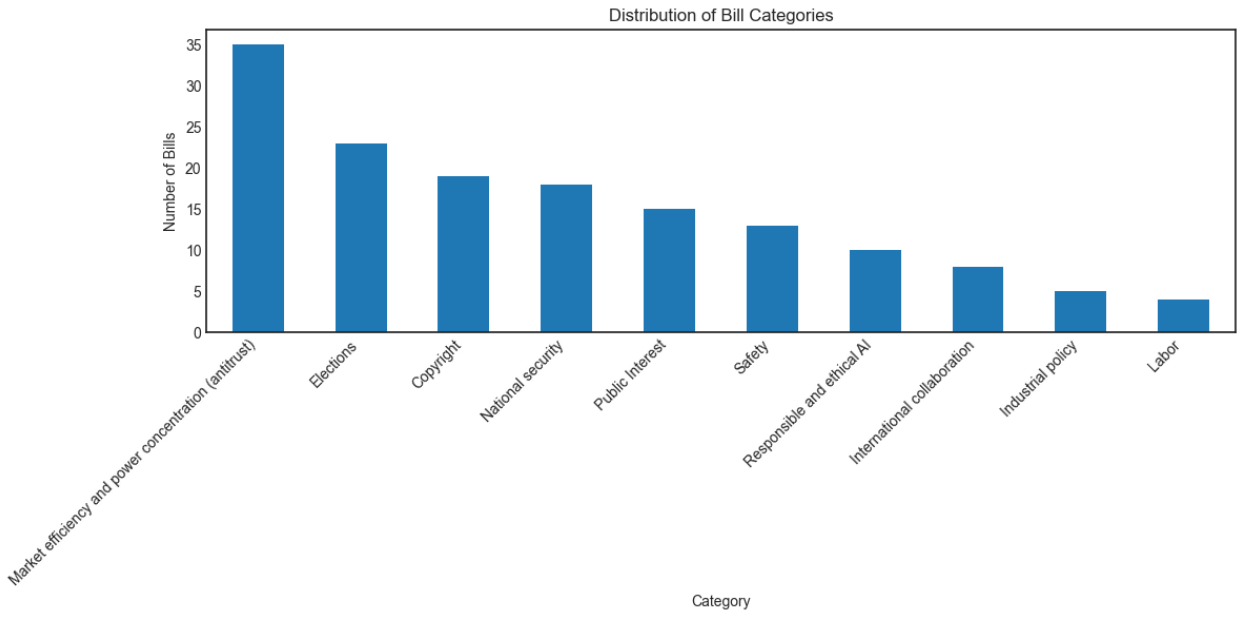

Figure 19 Distribution of Bill Categories (TF-IDF)

Latent Dirichlet Allocation (LDA) was then used to allow the corpus to suggest dominant policy themes within AI-related legislation using an unsupervised machine learning approach. We identified latent structures within legislative text, allowing us to classify bills based on their primary focus areas. Bill text was tokenized to remove stop words, punctuation, and special characters to be converted into vectorized representations following the same process and vectors as TF-IDF. This allowed the methods to ensure that frequently occurring words that are uninformative (e.g., "bill," "committee," "section") did not distort results. The model then extracted the most probable words per topic, developing its own similar topics.

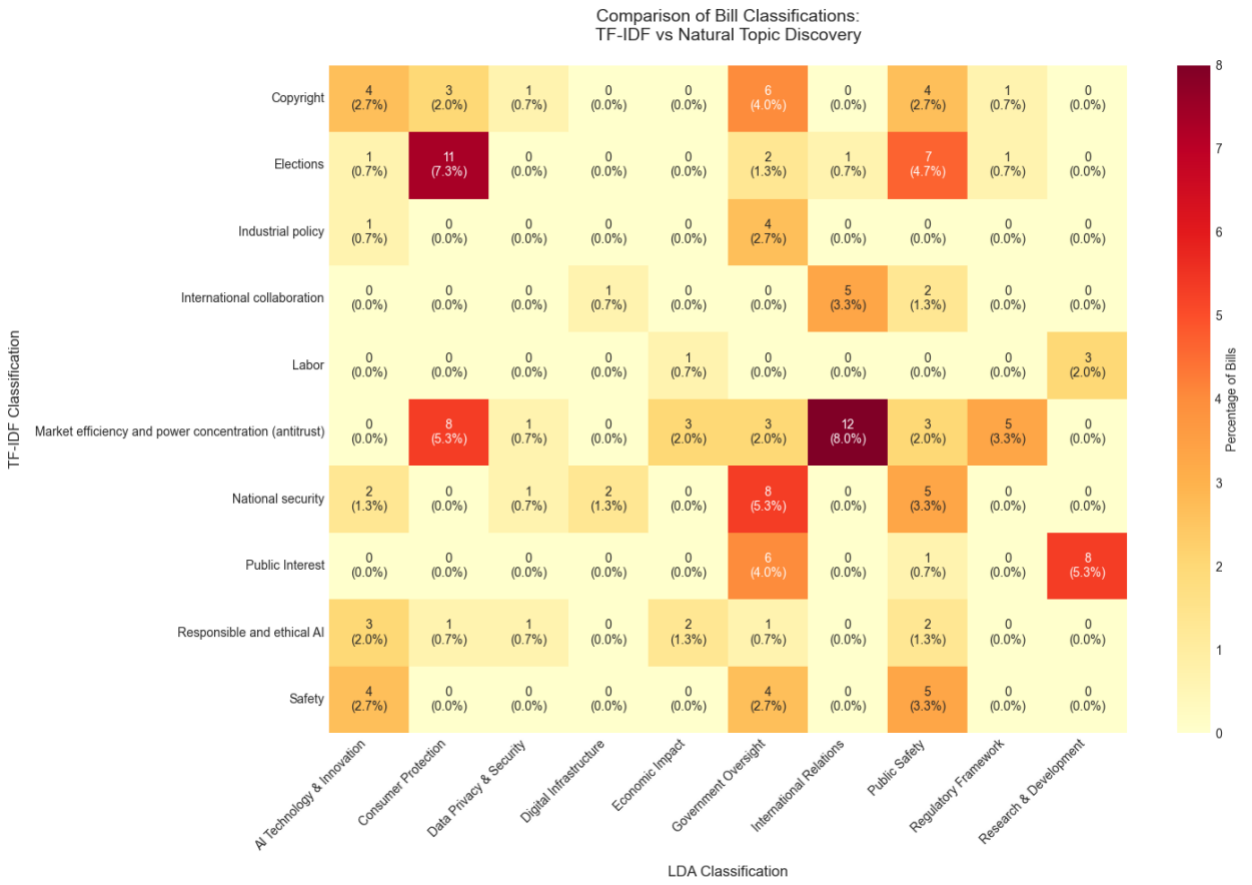

Figure 20 Comparison of Bill Classifications: TF-IDF vs Natural Topic Discovery

From a comparison of the proposed topics analyzed in TF-IDF and the topics examined in LDA, we observe that the most densely populated intersections show strong consistency across both techniques. For example, bills tagged as related to elections in TF-IDF overlap significantly with the proposed topics of government oversight and public safety, and bills tagged as market efficiency in TF-IDF overlap significantly with regulatory frameworks in LDA.

Validation with GPT

GPT has been demonstrated to perform on par and exceed the performance of Bidirectional Encoder Representations from Transformers (BERT) (Wang et al., 2024; Zhao et al., 2024) on classification and labeling tasks (Nguyen & Nguyen, 2024; Pelaez et al., 2024). We validated the proposed labels using GPT-4o's API using the same bill text processing techniques.

Sponsorship & Partisanship

To analyze sponsorship and partisanship Figure 24 and Figure 25, we extracted the processed sponsorship data using the Congress.gov API, identifying key attributes for each AI-related bill:

- Primary Sponsor Information:
  - Legislator name, party affiliation (Democratic, Republican, Independent), and chamber (House/Senate)
  - Whether the legislator had previously introduced AI-related legislation
- Cosponsor Information:
  - Number of cosponsors, including the same information as the primary sponsor
  - Whether cosponsors joined at the time of the bill's introduction or later in the legislative process
- Committee Assignments:
  - Committees responsible for reviewing each bill, including major AI-focused committees such as the
    - House Science, Space & Technology Committee
    - Senate Commerce, Science & Transportation Committee
    - House Judiciary Committee
    - Senate Armed Services Committee
- Bill Progression Tracking:
  - Legislative stage (e.g., introduced, passed committee, enacted into law).
  - Relationship between bipartisan sponsorship and legislative success rates.

Overall, Democrats introduced twice as many bills related to AI in the 118th Congress compared to Republicans, although the gap was more pronounced in the Senate and less pronounced in the House.

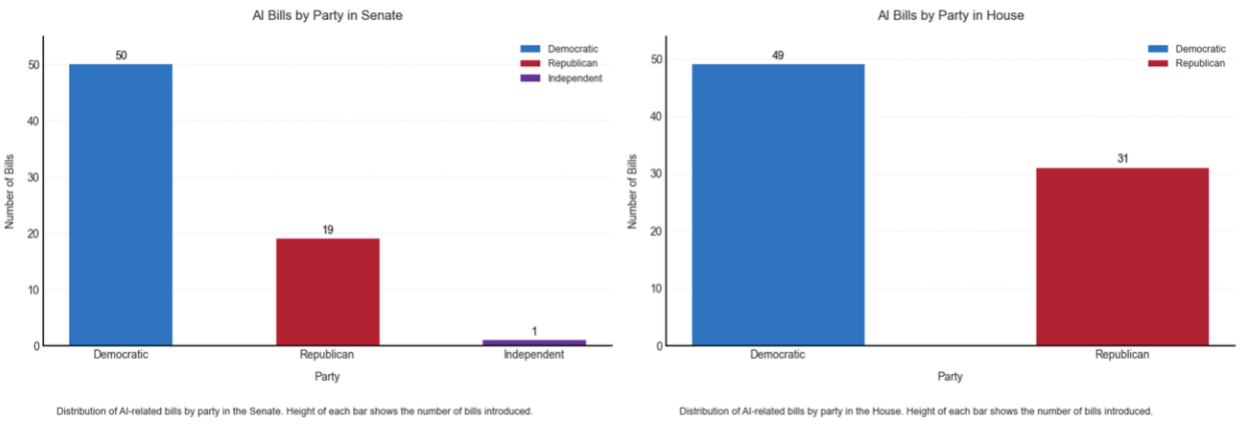

Figure 21 Distribution of party affiliation of primary bill sponsor

Committee Influence on AI Legislation

In Congress, the Senate Commerce, Science & Transportation Committee reviewed nearly half of all bills referred to committee in the Senate. However, the House was split more evenly, with a near-majority of bills going for review split between the House Energy and Commerce Committee and the House Science, Space, and Technology Committee. This is likely due to the split of commerce from the science committee in the House, whereas the two are merged in the Senate. See Figure 29 Most active committees.

Legislative Success

Although a large number of bills related to AI were introduced, only 11.3% made it out of committee, calculated by evaluating the progress data as defined in the data preprocessing step. See Figure 30 Bill progression.

Visualization

The data visualizations from Figure 22 to Figure 35 were generated by the previously discussed analyses.

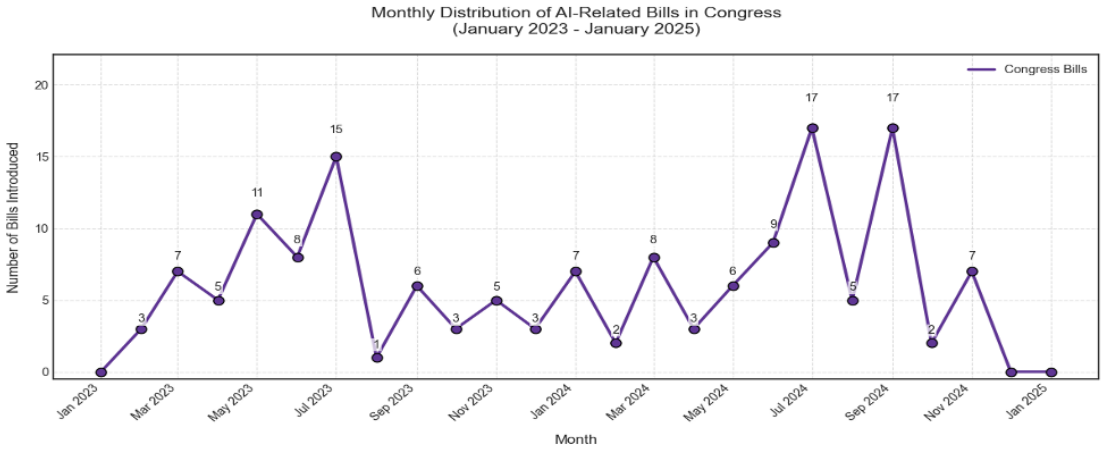

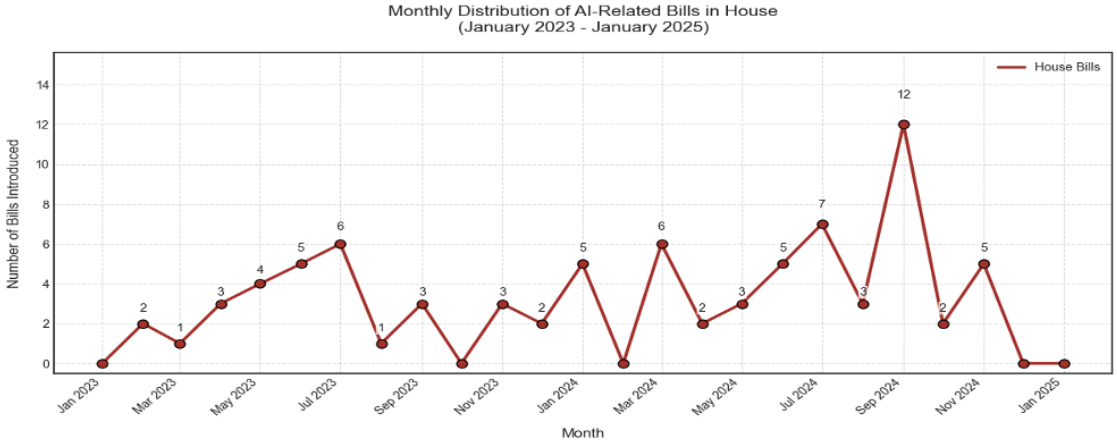

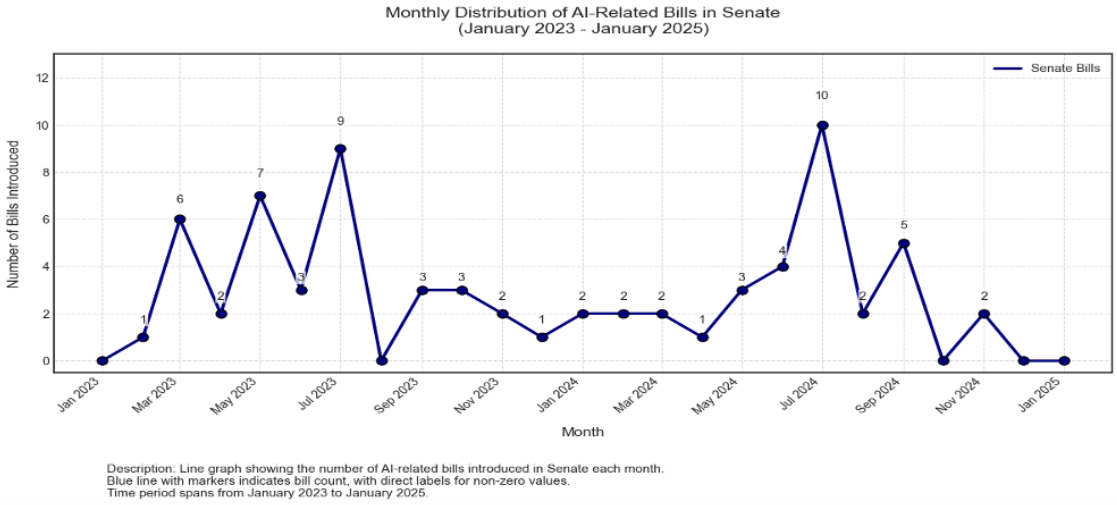

Figure 22 Timing of AI bill introduction

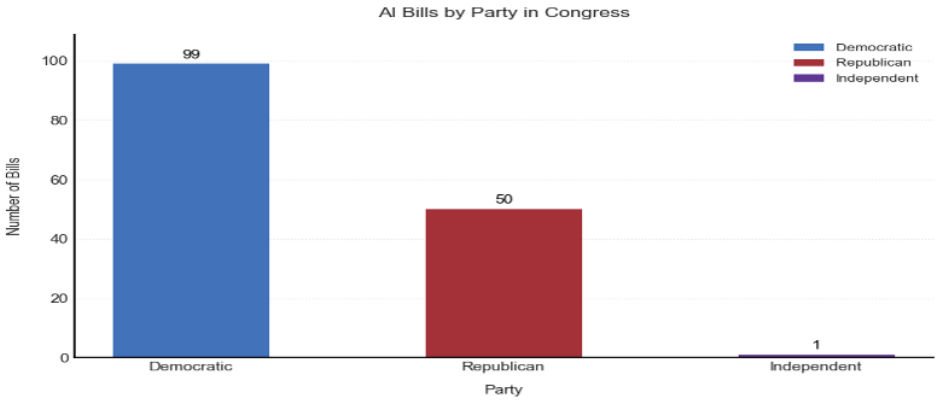

Distribution of AI-related bills by party in the Congress. Height of each bar shows the number of bills introduced.

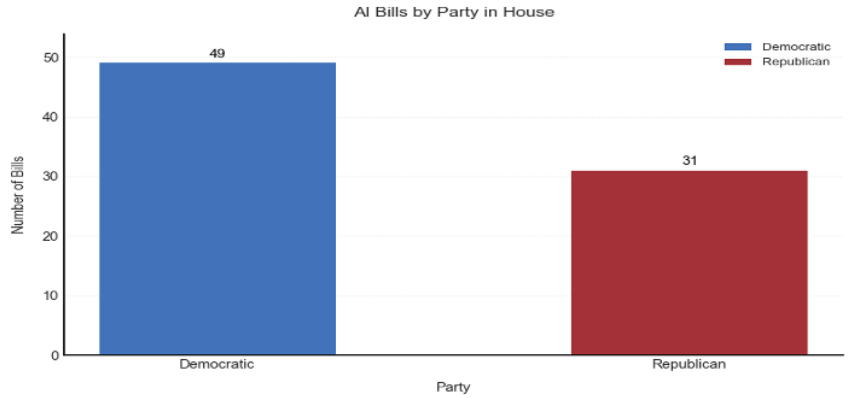

Distribution of AI-related bills by party in the House. Height of each bar shows the number of bills introduced.

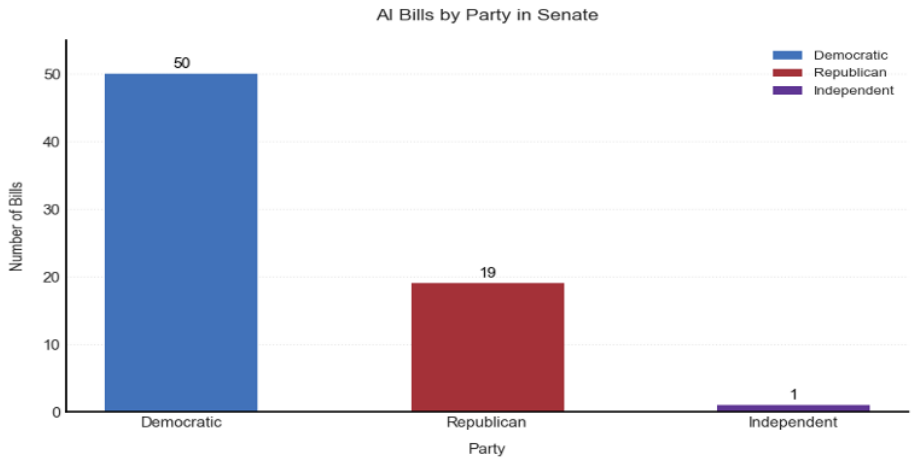

Distribution of AI-related bills by party in the Senate. Height of each bar shows the number of bills introduced.

Figure 23 AI bill sponsorship by chamber

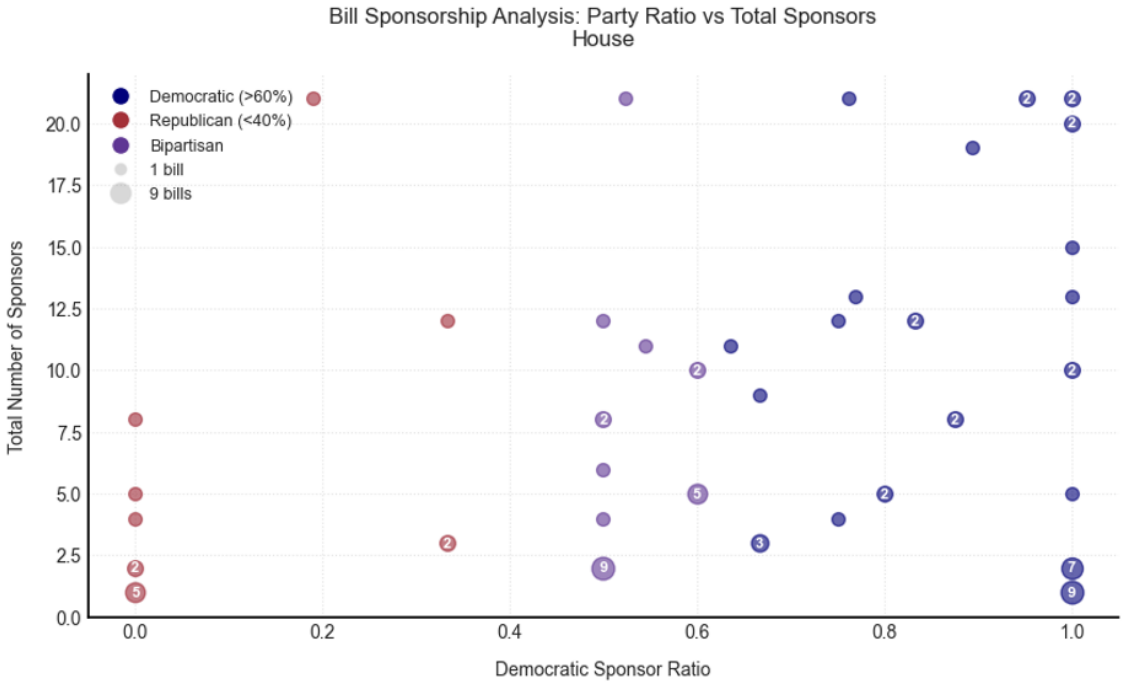

Scatter plot showing the relationship between Democratic sponsor ratio and total sponsors in the House.
Blue points indicate Democratic-led bills (>60% Democratic sponsors),
red points indicate Republican-led bills (<40% Democratic sponsors),
and purple points indicate bipartisan bills. Point size shows number of overlapping bills.

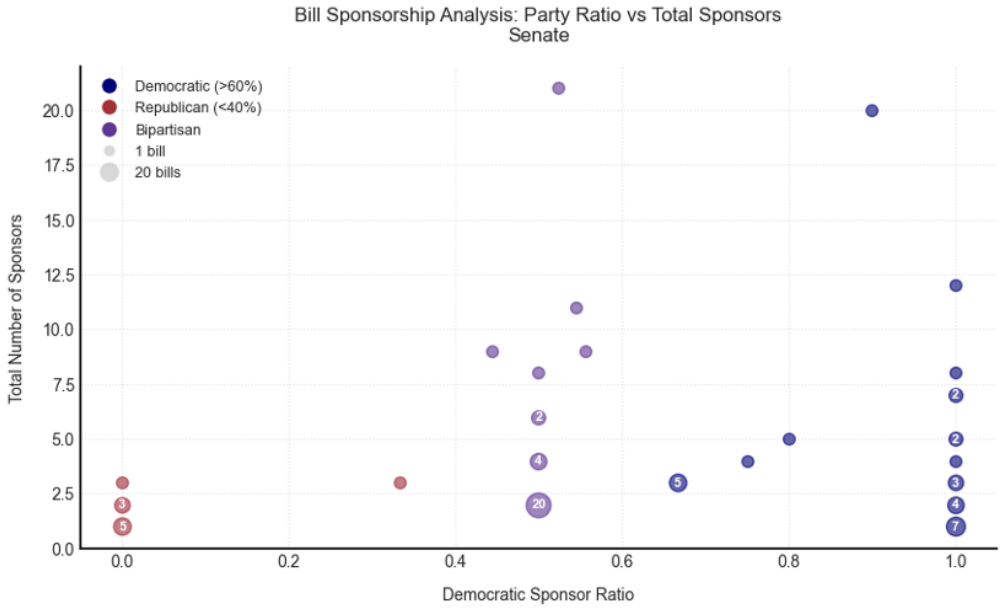

Scatter plot showing the relationship between Democratic sponsor ratio and total sponsors in the Senate.
Blue points indicate Democratic-led bills (>60% Democratic sponsors),
red points indicate Republican-led bills (<40% Democratic sponsors),
and purple points indicate bipartisan bills. Point size shows number of overlapping bills.

Figure 24 Bipartisanship by chamber of Congress

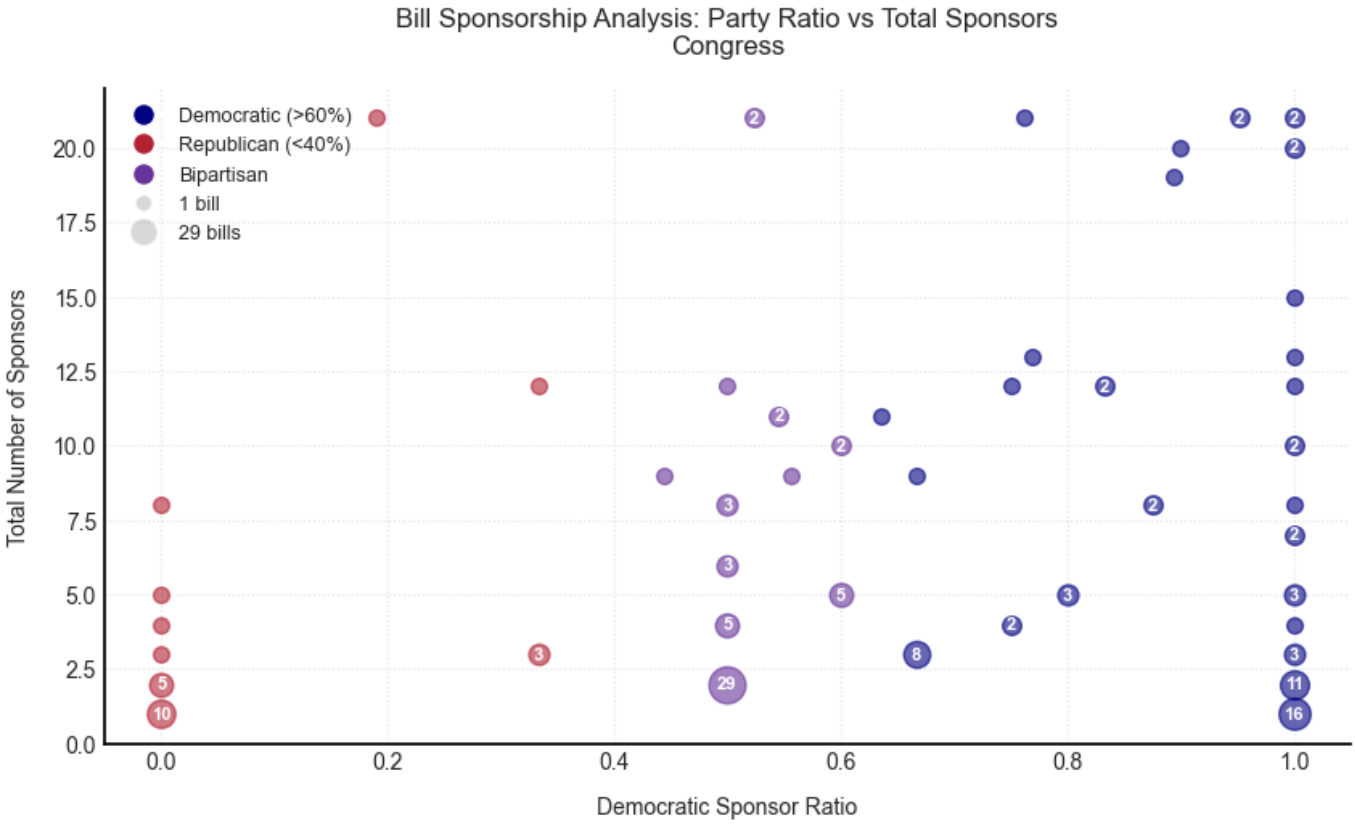

Scatter plot showing the relationship between Democratic sponsor ratio and total sponsors in the Congress.
Blue points indicate Democratic-led bills (>60% Democratic sponsors),
red points indicate Republican-led bills (<40% Democratic sponsors),
and purple points indicate bipartisan bills. Point size shows number of overlapping bills.

Figure 25 Bipartisanship analysis - Congress total

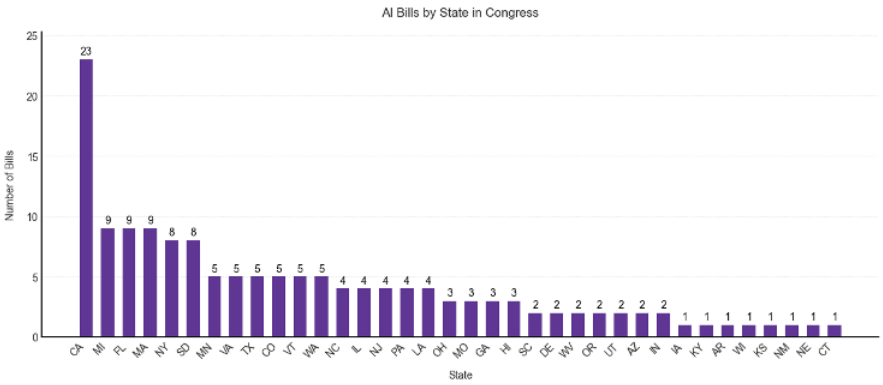

Distribution of AI-related bills by state in the Congress. Height of each bar shows the number of bills introduced across all 36 states.

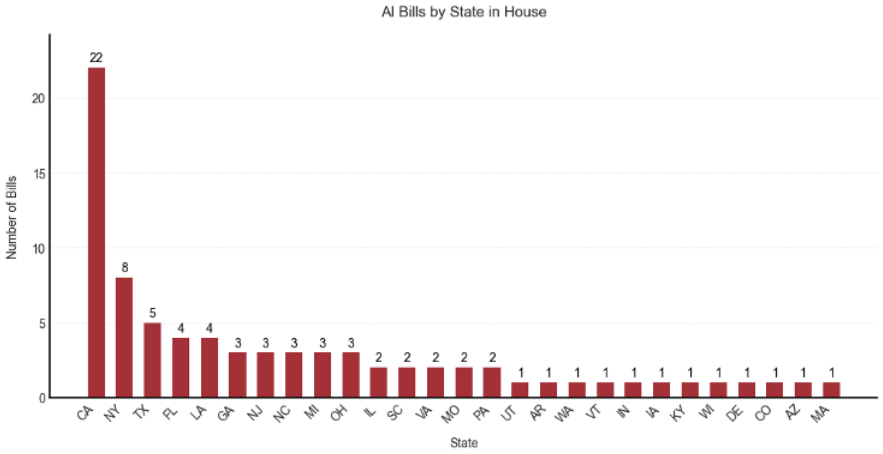

Distribution of AI-related bills by state in the House. Height of each bar shows the number of bills introduced across all 27 states.

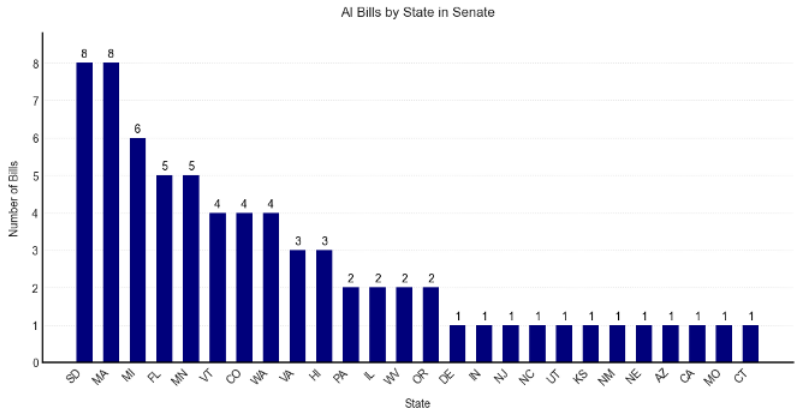

Distribution of AI-related bills by state in the Senate. Height of each bar shows the number of bills introduced across all 26 states.

Figure 26 AI bill sponsorship by state

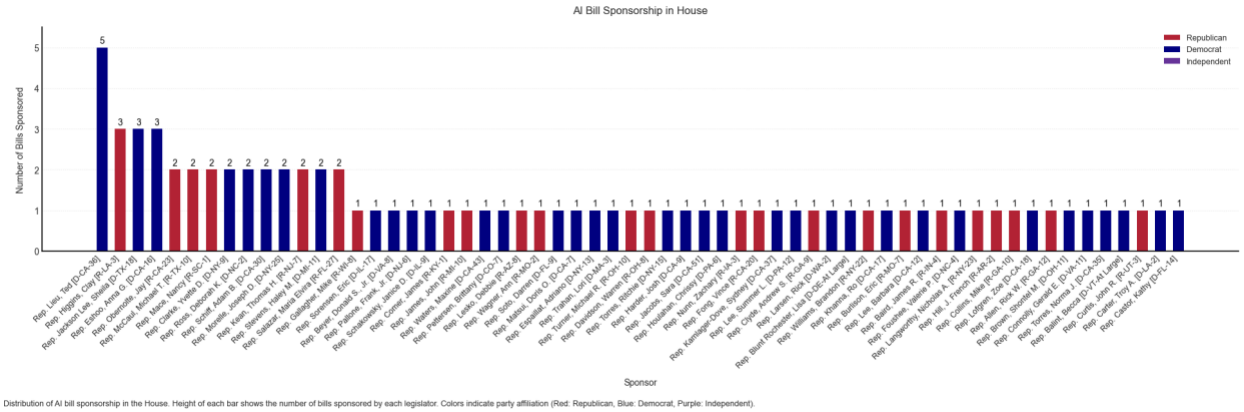

Figure 27 Most active members (House of Representatives)

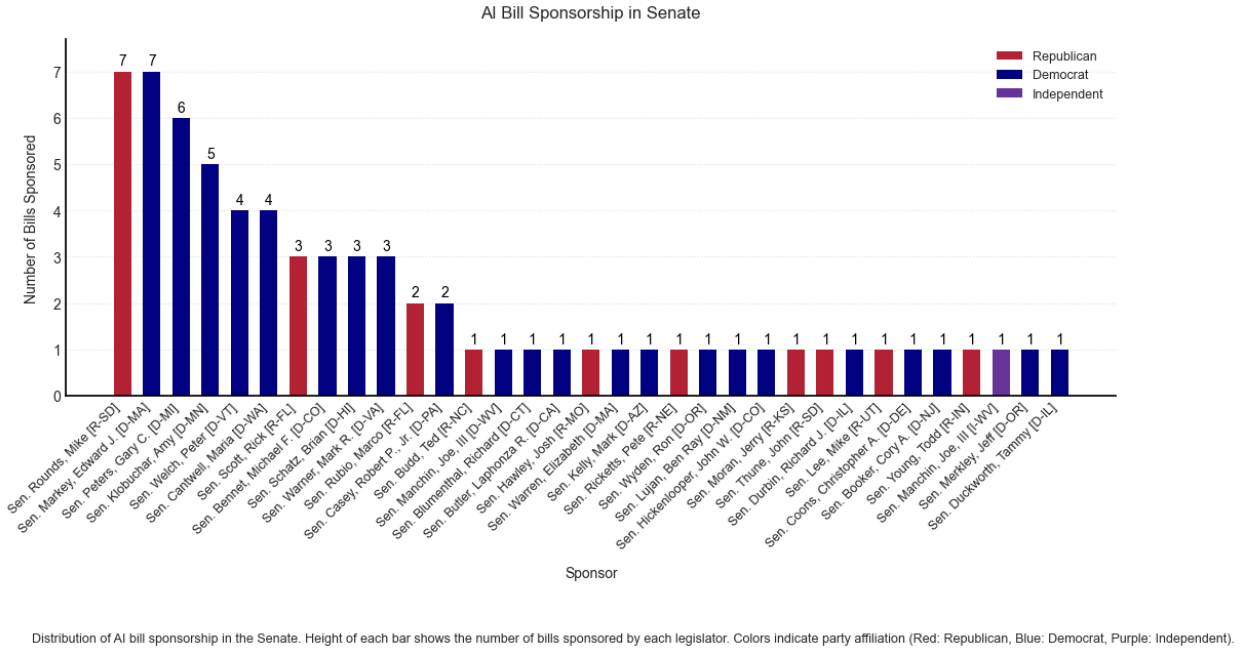

Figure 28 Most active members (Senate)

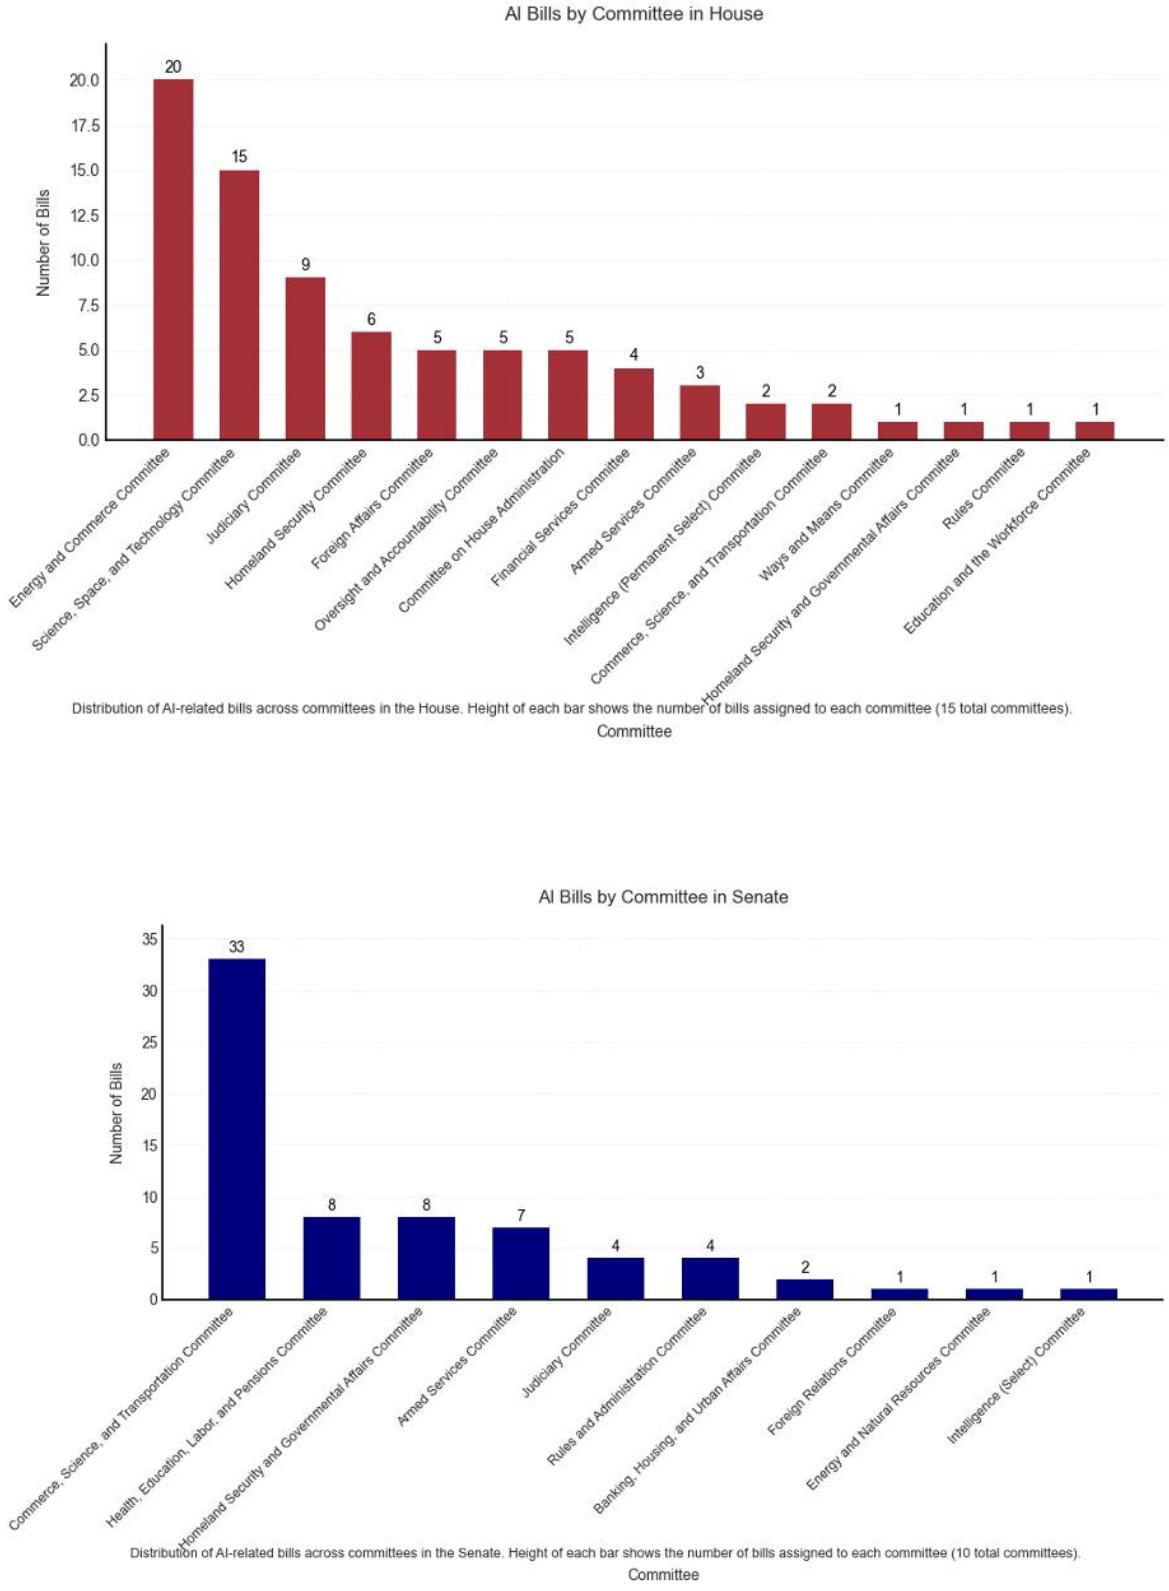

Figure 29 Most active committees

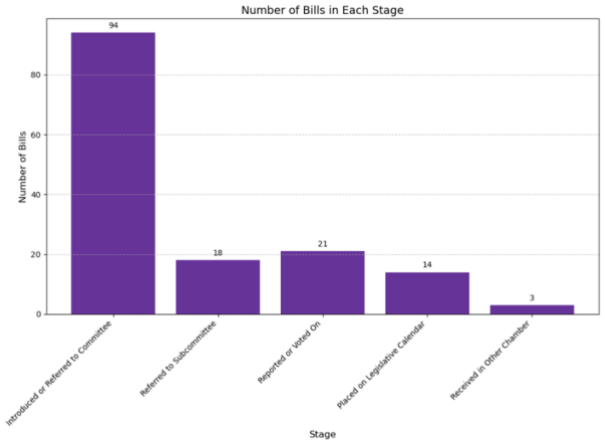

Figure 30 Bill progression

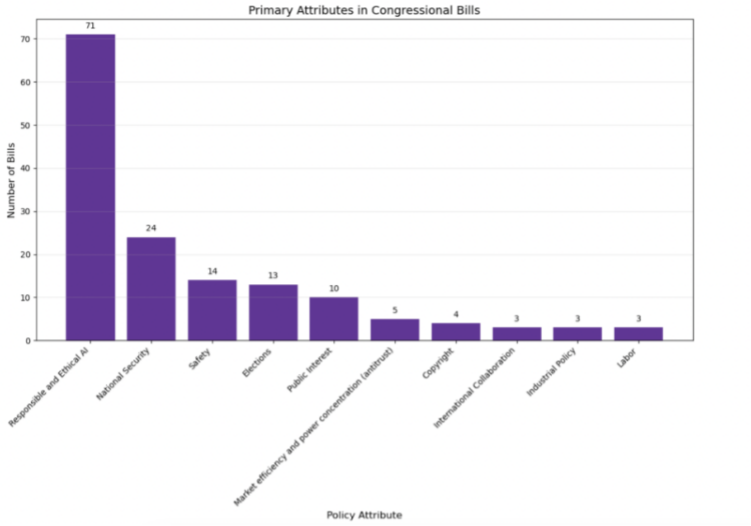

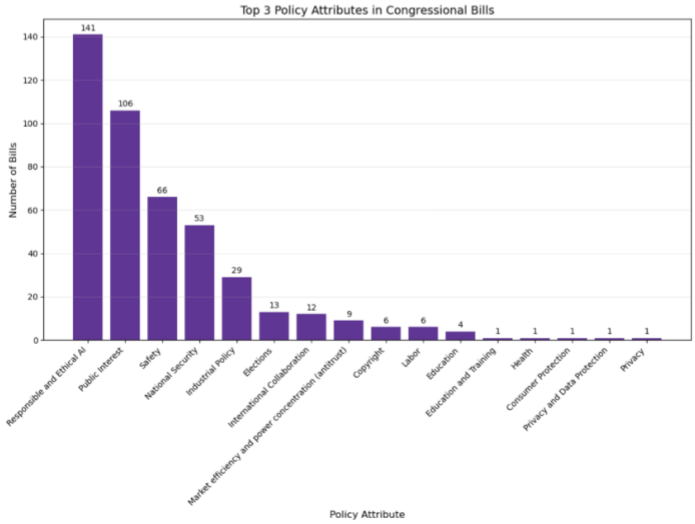

Figure 31 Bill policy attributes

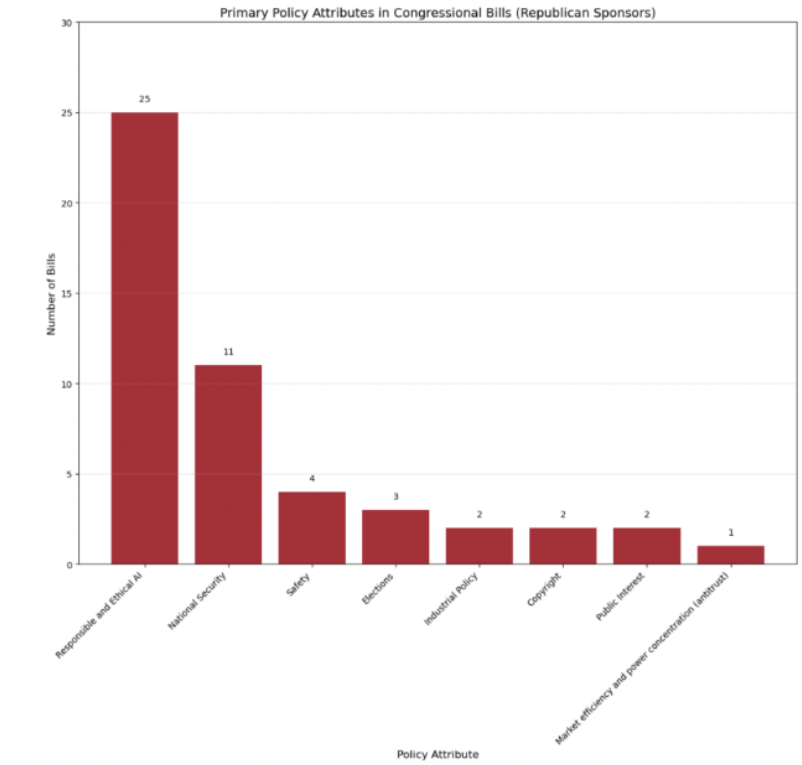

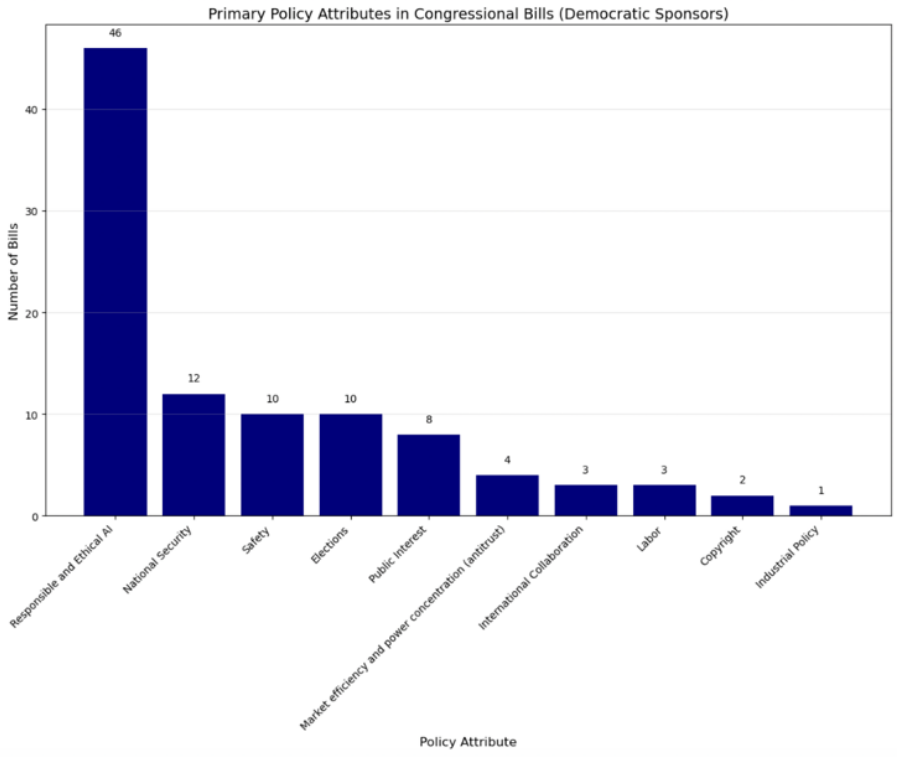

Figure 32 Primary bill policy attribute by party

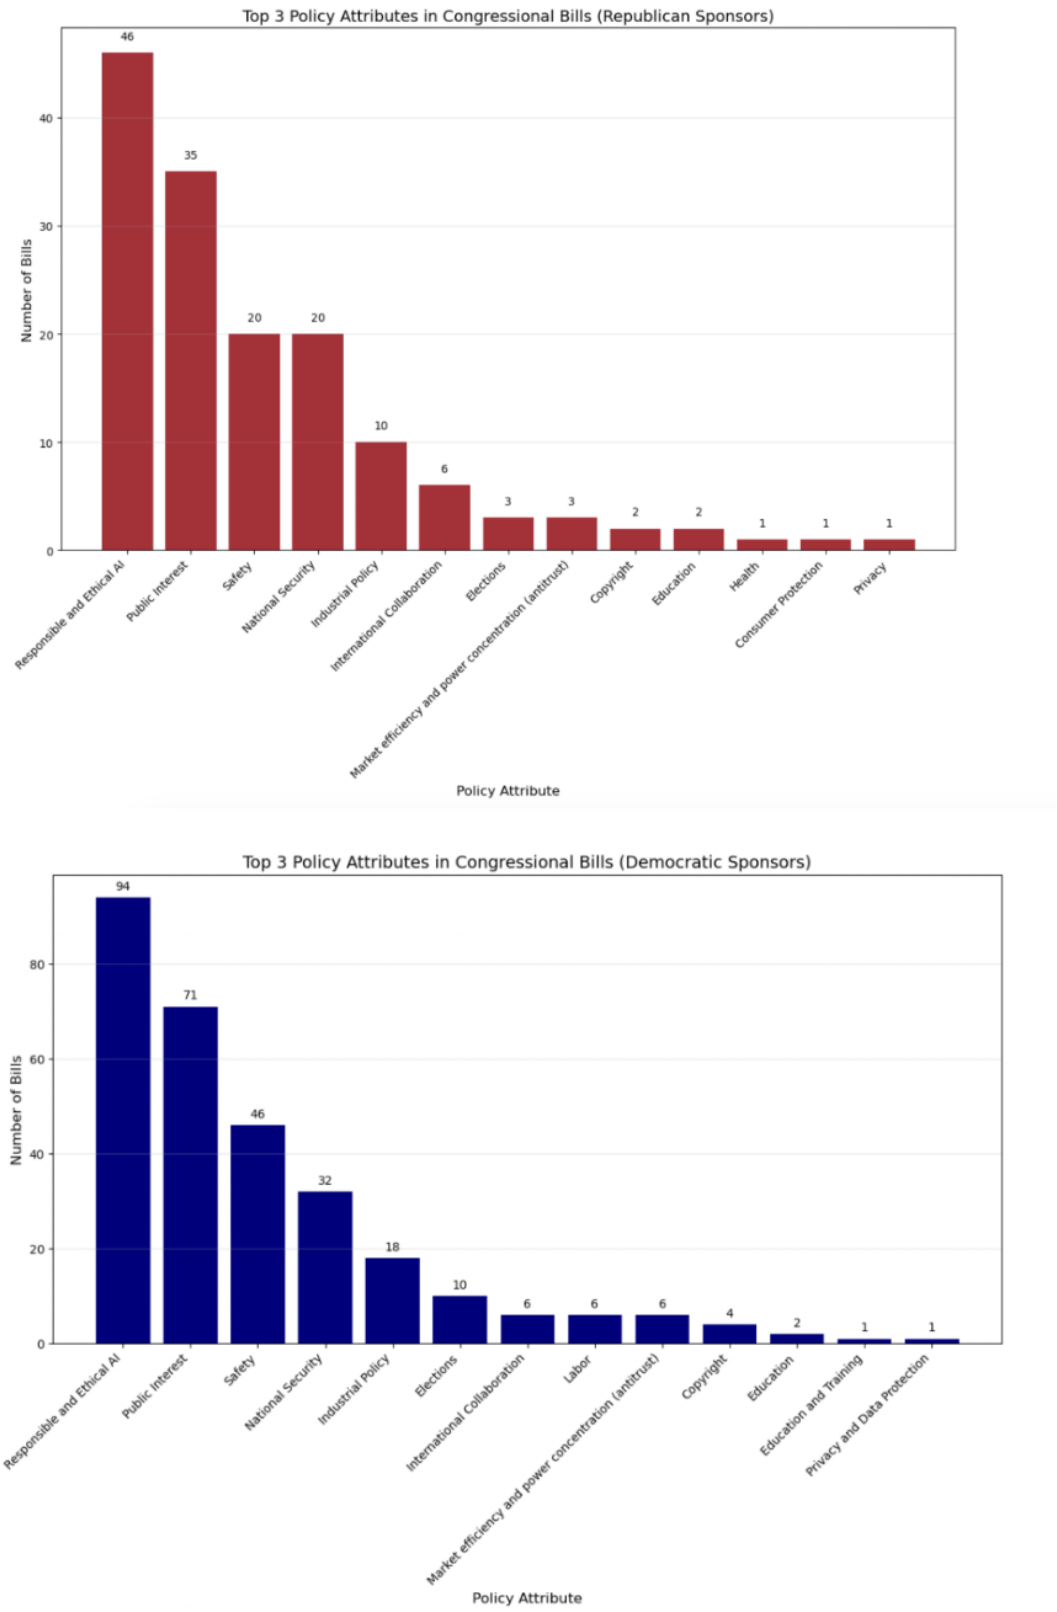

Figure 33 Top 3 bill policy attributes by party

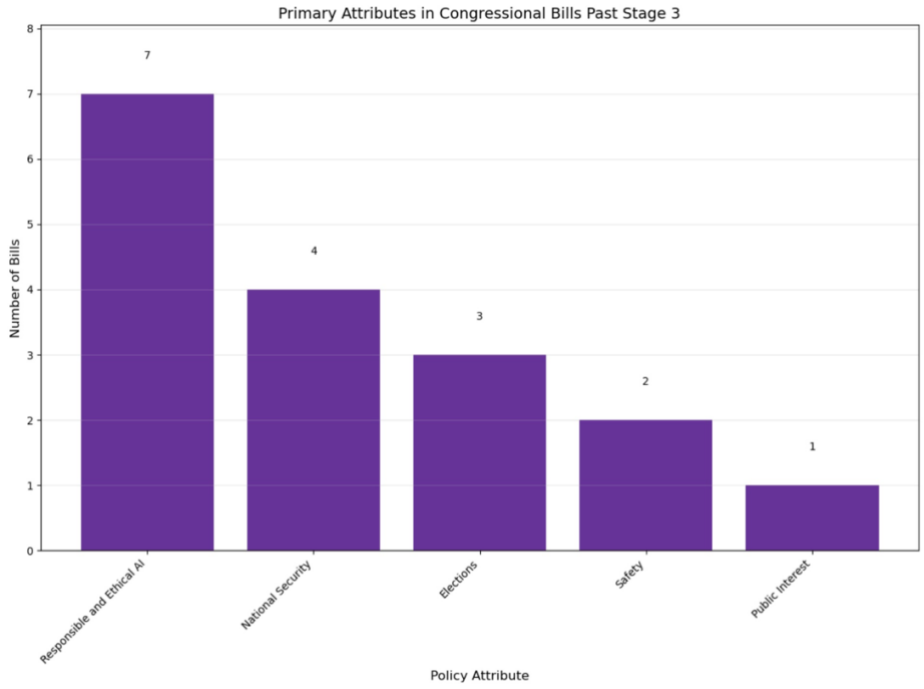

Figure 34 Primary policy attributes for the 17 bills placed on the legislative calendar (most progress)

Bill Number	Stage	Title
S.1564	4	AI Leadership Training Act
S.1865	4	TAG Act
S.2103	4	Intelligence Authorization Act for Fiscal Year 2024
S.2293	4	AI LEAD Act
S.2770	4	Protect Elections from Deceptive AI Act
H.R.3369	4	Artificial Intelligence Accountability Act
S.3897	4	Preparing Election Administrators for AI Act
H.R.3932	4	Intelligence Authorization Act for Fiscal Year 2024
S.4394	4	NSF AI Education Act of 2024
S.4487	4	Small Business Artificial Intelligence Training Act of 2024
S.4664	4	Department of Energy AI Act
S.4853	4	Ending FCC Meddling in Our Elections Act
H.R.1153	4	DATA Act
S.1409	4	Kids Online Safety Act
H.R.4755	7	Privacy Enhancing Technology Research Act
H.R.4814	7	Consumer Safety Technology Act
S.3696	7	DEFIANCE Act of 2024

Figure 35 Bills placed on the legislative calendar during the 118th Congress (total = 17)

References

Abdin, M., Aneja, J., Behl, H., Bubeck, S., Eldan, R., Gunasekar, S., Harrison, M., Hewett, R. J., Javaheripi, M., Kauffmann, P., Lee, J. R., Lee, Y. T., Li, Y., Liu, W., Mendes, C. C. T., Nguyen, A., Price, E., Rosa, G. de, Saarikivi, O., … Zhang, Y. (2024). *Phi-4 Technical Report* (No. arXiv:2412.08905). arXiv. https://doi.org/10.48550/arXiv.2412.08905

Abramov, M. (2024, July 5). *Regional and International AI Regulations and Laws in 2024*. Keymakr. https://keymakr.com/blog/regional-and-international-ai-regulations-and-laws-in-2024/

Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, *13*(1), 25–35. https://doi.org/10.1093/cjres/rsz022

ACLU. (2023, October 30). ACLU Statement on President Biden's Executive Order on Artificial Intelligence. *American Civil Liberties Union*. https://www.aclu.org/press-releases/aclu-statement-on-president-bidens-executive-order-on-artificial-intelligence

AI Policy Institute. (2023, October 30). *Vast Majority of US voters of All Political Affiliations Support President Biden's Executive Order on AI - AI Policy Institute*. https://theaipi.org/poll-biden-ai-executive-order-10-30/

Aibin, M., & Simic, M. (2024, July 8). *Energy Consumption of ChatGPT Responses | Baeldung on Computer Science*. https://www.baeldung.com/cs/chatgpt-large-language-models-power-consumption

Alms, N. (2024, November 8). *Trump promised to repeal Biden's AI executive order—Here's what to expect next*. Nextgov.Com. https://www.nextgov.com/artificial-intelligence/2024/11/trump-promised-repeal-bidens-ai-executive-order-heres-what-expect-next/400934/

Amazon. (2024, October 16). *Amazon signs agreements for innovative nuclear energy projects to address growing energy demands*. https://www.aboutamazon.com/news/sustainability/amazon-nuclear-small-modular-reactor-net-carbon-zero

American Economic Liberties Project. (n.d.). *Big Tech on Trial*. American Economic Liberties Project. Retrieved January 5, 2025, from https://www.economicliberties.us/tech-lawsuit-timelines/

Amodei, D., & Pottinger, M. (2025, January 6). *Opinion | Trump Can Keep America's AI Advantage*. WSJ. https://www.wsj.com/opinion/trump-can-keep-americas-ai-advantage-china-chips-data-eccdce91

Andreessen, M. (2023). The Techno-Optimist Manifesto. *CE Think Tank Newswire*.

Anthropic. (2024, October 31). *The case for targeted regulation*. https://www.anthropic.com/news/the-case-for-targeted-regulation

Ars Technica. (2011, November 15). *The 40th birthday of—Maybe—The first microprocessor, the Intel 4004*. Ars Technica. https://arstechnica.com/information-technology/2011/11/the-40th-birthday-ofmaybethe-first-microprocessor/

Ashford, N. A., & Hall, R. P. (2011). The Importance of Regulation-Induced Innovation for Sustainable Development. *Sustainability*, *3*(1), 270–292. https://doi.org/10.3390/su3010270

Ashford, N. A., & Heaton, G. R. (1976). Environmental and safety regulations: Reasons for their adoption and possible effects on technological innovation. *Environmental Policy and Law*, *1*(4), 172–176. https://doi.org/10.1016/S0378-777X(76)80123-0

ASHFORD, N., AYERS, C., & STONE, R. (1985). USING REGULATION TO CHANGE THE MARKET FOR INNOVATION. *The Harvard Environmental Law Review : HELR*, *9*(2), 419–466.

Atkinson, R. D. (2022). *How the IT Sector Powers the US Economy*. https://itif.org/publications/2022/09/19/how-the-it-sector-powers-the-us-economy/

Bachman, J. (2025, January 3). *FCC net neutrality rule among first to fall in Loper Bright's aftermath*. Legal Dive. https://www.legaldive.com/news/fcc-net-neutrality-rule-among-first-to-fall-in-loper-brights-aftermath/736429/

Baer, B. (2025, January 7). What the US Learned and the EU Should Consider About National Champions. *ProMarket*. http://www.promarket.org/2025/01/07/what-the-us-learned-and-the-eu-should-consider-about-national-champions/

Barroso, L. R. (2024, December 16). *Artificial Intelligence: Promises, Risks, and Regulation*. https://www.hks.harvard.edu/centers/carr/publications/artificial-intelligence-promises-risks-and-regulation

Bennett Moses, L. (2007). Recurring Dilemmas: The Law's Race to Keep Up With Technological Change. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.979861

Bergmann, D. (2024, April 5). *What is mixture of experts?* | *IBM*. https://www.ibm.com/think/topics/mixture-of-experts

Bick, A., Blandin, A., & Deming, D. J. (2024). *The Rapid Adoption of Generative AI* (Working Paper No. 32966). National Bureau of Economic Research. https://doi.org/10.3386/w32966

Blaszczyk, M., McGovern, G., & Stanley, K. D. (2024). *Artificial Intelligence Impacts on Copyright Law*. RAND Corporation. https://www.rand.org/pubs/perspectives/PEA3243-1.html

Braun, M., Vellery, A., & Benizri, I. (2024, October 23). *Measures in Support of Innovation in the European Union's AI Act – AI Regulatory Sandboxes*. https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/20241023-measures-in-support-of-innovation-in-the-european-unions-ai-act-ai-regulatory-sandboxes

Brennan Center for Justice. (2024, December 2). *Artificial Intelligence Legislation Tracker* | *Brennan Center for Justice*. https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker

Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies “Engines of growth”? *Journal of Econometrics*, *65*(1), 83–108. https://doi.org/10.1016/0304-4076(94)01598-T

Broad, W. J. (1987, May 12). DOES THE FEAR OF LITIGATION DAMPEN THE DRIVE TO INNOVATE? *The New York Times*. https://www.nytimes.com/1987/05/12/science/does-the-fear-of-litigation-dampen-the-drive-to-innovate.html

Buchanan, B. (2020, August). The AI Triad and What It Means for National Security Strategy. *Center for Security and Emerging Technology*. https://cset.georgetown.edu/publication/the-ai-triad-and-what-it-means-for-national-security-strategy/

Carvao, P. (2023, May 11). *Resetting the Rules for Tech to Preserve the Public Interest* | *TechPolicy.Press*. Tech Policy Press. https://techpolicy.press/resetting-the-rules-for-tech-to-preserve-the-public-interest

Carvao, P. (2024a). 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. (2024b, October 26). AI and the Environment [Substack newsletter]. *Tech and Democracy*. https://carvao.substack.com/p/ai-and-the-environment

Carvao, P. (2025, January 14). Sounding the alarm in AI and National Security [Substack newsletter]. *Tech and Democracy*. https://carvao.substack.com/p/ai-natsec

Carvao, P., & Ancheva, S. (2024, October 4). California, Chevron Doctrine, and the State of AI Policy in the United States [Substack newsletter]. *Tech and Democracy*. https://carvao.substack.com/p/california-chevron-doctrine-and-the

CBRE. (2024, June 24). *Global Data Center Trends 2024*. https://www.cbre.com/insights/reports/global-data-center-trends-2024

Chan, A., Salganik, R., Markelius, A., Pang, C., Rajkumar, N., Krasheninnikov, D., Langosco, L., He, Z., Duan, Y., Carroll, M., Lin, M., Mayhew, A., Collins, K., Molamohammadi, M., Burden, J., Zhao, W., Rismani, S., Voudouris, K., Bhatt, U., … Maharaj, T. (2023). *Harms from Increasingly Agentic Algorithmic Systems*. 651–666. https://doi.org/10.1145/3593013.3594033

ChatGPT Revenue and Usage Statistics (2024). (n.d.). Business of Apps. Retrieved December 10, 2024, from https://www.businessofapps.com/data/chatgpt-statistics/

Cheung, S. (2024, March 21). *Central Washington Ranks Eighth, Seattle Tenth for North American Data Center Leasing in 2023*. https://www.cbre.com/press-releases/central-washington-ranks-eighth-seattle-tenth-for-north-american-data-center-leasing-in-2023

Chollet, F. (2024, December 20). *OpenAI o3 Breakthrough High Score on ARC-AGI-Pub*. ARC Prize. https://arcprize.org/blog/oai-o3-pub-breakthrough

Clark, J., & Hadfield, G. K. (2019). Regulatory Markets for AI Safety. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2001.00078

Codd, E. F. (1970). A relational model of data for large shared data banks. *Commun. ACM*, *13*(6), 377–387. https://doi.org/10.1145/362384.362685

Cong, W., & Yeping, Y. (2023, October 18). *China launches Global AI Governance Initiative, offering an open approach in contrast to US blockade—Global Times*. https://www.globaltimes.cn/page/202310/1300092.shtml

Congress.gov. (n.d.). *Browse U.S. Legislative Information by Policy Area—119th Congress (2025-2026)* [Legislation]. Retrieved February 5, 2025, from https://www.congress.gov/browse/policyarea

Czarnecka, K. (2024, May 8). AI's carbon footprint—How does the popularity of artificial intelligence affect the climate? *Plan Be Eco*. https://planbe.eco/en/blog/ais-carbon-footprint-how-does-the-popularity-of-artificial-intelligence-affect-the-climate/

Daum, J. (2023, April 14). Overview of Draft Measures on Generative AI. *China Law Translate*. https://www.chinalawtranslate.com/overview-of-draft-measures-on-generative-ai/

Davenport, C., Singer, B., Mehta, N., Lee, B., Mackay, J., Modak, A., Corbett, B., Miller, J., Hari, T., Ritchi, J., Delaney, M., Revich, J., Jaiya, J., Venugopal, V., Cash, N., & Halferty, O. (2024, April 29). *Generational Growth: AI, data centers and the coming US power demand surge*. https://www.goldmansachs.com/insights/goldman-sachs-research/generational-growth-ai-data-centers-and-the-coming-us-power-demand-surge

Dean, J. (2023, October 24). Technology (A Special Report)—Will AI Bring Utopia or Dystopia? A Venture Capitalist Picks Utopia: Vinod Khosla says artificial intelligence will cause great abundance—More than enough to share. *Wall Street Journal*. ProQuest Central; ProQuest One Business; The Wall Street Journal. http://search.proquest.com.ezp-prod1.hul.harvard.edu/newspapers/technology-special-report-will-ai-bring-utopia/docview/2880628171/se-2?accountid=11311

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4573321

Eccles, R. G. (2024, August 31). *Microsoft Can Take The Lead In Small Modular Reactors For Powering AI*. Forbes. https://www.forbes.com/sites/bobeccles/2024/08/31/microsoft-can-take-the-lead-in-small-modular-reactors-for-powering-ai/

Eggers, W., Walsh, S., Joergensen, C., & Kishnani, P. (2023, March 23). *Regulation that enables innovation*. Deloitte Insights.

https://www2.deloitte.com/us/en/insights/industry/public-sector/government-trends/2023/regulatory-agencies-and-innovation.html

Encyclopedia of Artificial Intelligence: The Past, Present, and Future of AI. (2021). *ProtoView*.

Epoch AI. (2024, June 19). *Data on Notable AI Models*. Epoch AI. https://epoch.ai/data/notable-ai-models

European Union. (n.d.). *CE marking – obtaining the certificate, EU requirements*. Your Europe. Retrieved January 4, 2025, from https://europa.eu/youreurope/business/product-requirements/labels-markings/ce-marking/index_en.htm

Finkle, N. (2025, January 13). NVIDIA Statement on the Biden Administration's Misguided "AI Diffusion" Rule. *NVIDIA Blog*. https://blogs.nvidia.com/blog/ai-policy/

Friedman, L. M. (Lawrence M. (2023). *Law, science, and technology: Historical and social context*. Rowman & Littlefield, an imprint of The Roman & Littlefield Publishing Group, Inc.

Future of Life Institute. (n.d.). *Article 57: AI Regulatory Sandboxes | EU Artificial Intelligence Act*. Retrieved January 9, 2025, from https://artificialintelligenceact.eu/article/57/

Ghaffary, S. (2024, September 24). *OpenAI Pitched White House on Unprecedented Data Center Buildout*. Yahoo Finance. https://finance.yahoo.com/news/openai-pitched-white-house-unprecedented-000550020.html

Goldman Sachs. (2024, May 14). *AI is poised to drive 160% increase in data center power demand*. https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand

Gornet, M., & Maxwell, W. (2024). The European approach to regulating AI through technical standards. *Internet Policy Review*, *13*(3), 1–27. https://doi.org/10.14763/2024.3.1784

GovTrack.us. (2025, February). *Historical Statistics about Legislation in the U.S. Congress -*. GovTrack.Us. https://www.govtrack.us/congress/bills/statistics

Green, A., Tai, H., Noffsinger, J., Sachdeva, P., Bhan, A., & Sharma, R. (2024, September 17). *Data centers and AI: How the energy sector can meet power demand | McKinsey*. https://www.mckinsey.com/industries/private-capital/our-insights/how-data-centers-and-the-energy-sector-can-sate-ais-hunger-for-power

Green-Lowe, J. (2024, May 16). *The Senate's AI Roadmap to Nowhere | Center for AI Policy | CAIP*. Center for AI Policy. https://www.centeraipolicy.org/work/the-senates-ai-roadmap-to-nowhere

Gutiérrez, J. D. (2024, August 16). *Consultation paper on AI regulation: Emerging approaches across the world—UNESCO Digital Library*. https://unesdoc.unesco.org/ark:/48223/pf0000390979

Harari, Y. N. (2024). *Nexus: A brief history of information networks from the Stone Age to AI* (First U.S. edition.). Random House.

Harris, L. (2023). Artificial Intelligence: Overview, Recent Advances, and Considerations for the 118th Congress. CRS Report. *Congressional Research Service*.

Hendrix, J., Miller, G., & Lennett, B. (2024, May 15). *US Senate AI Working Group Releases Policy Roadmap | TechPolicy.Press*. Tech Policy Press. https://techpolicy.press/us-senate-ai-working-group-releases-policy-roadmap

Herdman, R. C. (1993). *Technologies Underlying Weapons of Mass Destruction*.

Hochreiter, S., & Schmidhuber, Jr. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Houlihan Lockey. (2024, Summer). *Real Estate Market Update*. Houlihan Lokey. https://hl.com/insights/real-estate-market-update/

House Committee on Science, Space, and Technology. (2024, December 17). *House Bipartisan Task Force on Artificial Intelligence Delivers Report*. House Committee on Science Space & Tech - Republicans. https://science.house.gov/2024/12/house-bipartisan-task-force-on-artificial-intelligence-delivers-report

Hurst, J. W. (2022). CHAPTER EIGHT TECHNOLOGY AND THE LAW: THE AUTOMOBILE. *Wisconsin Law Review*, *3*, 463–531.

Interesse, G. (2024, May 30). *China Releases New Draft Regulations for Generative AI*. China Briefing News. https://www.china-briefing.com/news/china-releases-new-draft-regulations-on-generative-ai/

ITI. (2025, January 13). *ITI: New Export Controls Threaten to Discourage the Use of U.S. Technology - Information Technology Industry Council*. https://www.itic.org/news-events/news-releases/iti-new-export-controls-threaten-to-discourage-the-use-of-u-s-technology

Iyer, P., & Hendrix, J. (2024, December 20). *Reactions to the Bipartisan US House AI Task Force Report | TechPolicy.Press*. Tech Policy Press. https://techpolicy.press/reactions-to-the-bipartisan-us-house-ai-task-force-report

Kagan, J. (2023, January 30). *Underwriters Laboratories (UL): Meaning, Overview, History*. Investopedia. https://www.investopedia.com/terms/u/underwriters-laboratories-ul.asp

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., & Amodei, D. (2020). *Scaling Laws for Neural Language Models* (No. arXiv:2001.08361). arXiv. https://doi.org/10.48550/arXiv.2001.08361

Krishan, N. (2023, November 9). Tech groups push back on Biden AI executive order, raising concerns that it could crush innovation. *FedScoop*. https://fedscoop.com/tech-groups-push-back-on-biden-ai-executive-order-raising-concerns-that-it-could-crush-innovation/

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, *25*.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, *1*(4), 541–551. https://doi.org/10.1162/neco.1989.1.4.541

Library of Congress. (n.d.). *Congress.gov | Library of Congress* [Legislation]. Retrieved January 5, 2025, from https://www.congress.gov/

Liou, J. (2023, September 13). *What are Small Modular Reactors (SMRs)?* [Text]. IAEA. https://www.iaea.org/newscenter/news/what-are-small-modular-reactors-smrs

Lynch, S. (2024, April 15). *AI Index: State of AI in 13 Charts*. https://hai.stanford.edu/news/ai-index-state-ai-13-charts

MacCarthy, M. (2023). *Regulating digital industries: How public oversight can encourage competition, protect privacy, and ensure free speech*. Brookings Institution Press.

Mandel, G. N. (2007). *History Lessons for a General Theory of Law and Technology*.

Marcus, G. (2024, November 9). CONFIRMED: LLMs have indeed reached a point of diminishing returns [Substack newsletter]. *Marcus on AI*. https://garymarcus.substack.com/p/confirmed-llms-have-indeed-reached

Mathieu, E. (2024, May 21). *The price of computer storage has fallen exponentially since the 1950s*. Our World in Data. https://ourworldindata.org/data-insights/the-price-of-computer-storage-has-fallen-exponentially-since-the-1950s

Matt O'Brien & Sarah Parvini. (2024, July 12). *US senators call out Big Tech's new approach to poaching talent, products from smaller AI startups*. AP News. https://apnews.com/article/ai-artificial-intelligence-acquihires-amazon-adept-wyden-fa3cd0502a757e5a9ccb83d00bf9ad55

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence: August 31, 1955. *AI Magazine*, *27*(4), 12–14. ProQuest Central; ProQuest One Business; Social Science Premium Collection.

Ministry of Foreign Affairs The People's Republic of China. (2024, September 30). *Promoting Development for All and Bridging the AI Divide_Ministry of Foreign Affairs of the*

People's Republic of China. https://www.mfa.gov.cn/eng/wjbzhd/202409/t20240930_11501255.html

Moore, G. E. (2006). Cramming more components onto integrated circuits, Reprinted from Electronics, volume 38, number 8, April 19, 1965, pp. 114 ff. *IEEE Solid-State Circuits Society Newsletter*, *11*(3), 33–35.

Multistate.ai. (n.d.). *Artificial Intelligence (AI) Legislation*. Multistate.Ai. Retrieved January 10, 2025, from https://www.multistate.ai/artificial-intelligence-ai-legislation

National Conference of State Legislatures. (2024, September 9). *Artificial Intelligence 2024 Legislation*. https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation

Newman, J., Dubal, V., Viljoen, S., Ajunwa, I., Guggenberger, N., Bietti, E., Jackson, J., & Tan, J. (2023, December 4). *Seven Reactions to Biden's Executive Order on Artificial Intelligence*. LPE Project. https://lpeproject.org/blog/seven-reactions-to-bidens-executive-order-on-artificial-intelligence/

Nguyen, C., & Nguyen, L.-M. (2024). *Employing Label Models on ChatGPT Answers Improves Legal Text Entailment Performance* (No. arXiv:2401.17897). arXiv. https://doi.org/10.48550/arXiv.2401.17897

Nihill, C. (2024, December 20). Bipartisan Senate bill would establish an AI safety office in Commerce. *FedScoop*. https://fedscoop.com/bipartisan-senate-bill-ai-safety-office-commerce/

NTIA. (2024a, March 27). *AI Accountability Policy Report | National Telecommunications and Information Administration*. https://www.ntia.gov/issues/artificial-intelligence/ai-accountability-policy-report

NTIA. (2024b, July 30). *Dual-Use Foundation Models with Widely Available Model Weights Report | National Telecommunications and Information Administration*.

https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report

NVIDIA Blackwell Architecture. (n.d.). NVIDIA. Retrieved December 10, 2024, from https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/

NYU's Center for Social Media and Politics. (2024, December 12). *The State of State Technology Policy: 2024 Report*. NYU's Center for Social Media and Politics. https://csmapnyu.org/impact/policy/the-state-of-state-technology-policy-2024-report

OECD. (2023, July 12). *Regulatory sandboxes in artificial intelligence*. OECD. https://www.oecd.org/en/publications/regulatory-sandboxes-in-artificial-intelligence_8f80a0e6-en.html

OECD. (2024). *Better regulation and innovation*. OECD. https://www.oecd.org/en/topics/better-regulation-and-innovation.html

Oklo Inc. (n.d.). *Oklo Inc. - Investors—Governance—Board of directors*. https://oklo.com/investors/governance/board-of-directors/default.aspx

OpenAI. (2024, September 12). *Learning to Reason with LLMs*. https://openai.com/index/learning-to-reason-with-llms/

OpenAI. (2025, January 21). *Announcing The Stargate Project*. https://openai.com/index/announcing-the-stargate-project/

Pelaez, S., Verma, G., Ribeiro, B., & Shapira, P. (2024). Large-scale text analysis using generative language models: A case study in discovering public value expressions in AI patents. *Quantitative Science Studies*, *5*(1), 153–169. https://doi.org/10.1162/qss_a_00285

Pichai, S., Hassabis, D., & Kavukcuoglu, K. (2024, December 11). *Introducing Gemini 2.0: Our new AI model for the agentic era*. Google. https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/

Plumer, B. (2024, September 20). Three Mile Island Plans to Reopen as Demand for Nuclear Power Grows. *The New York Times*. https://www.nytimes.com/2024/09/20/climate/three-mile-island-reopening.html

Pope, A. (2024, April 10). *NYT v. OpenAI: The Times's About-Face*. Harvard Law Review. https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/

PORTER M. (1991). America's Green Strategy. *Scientific American*, 86-.

Radford, A. (2018). *Improving language understanding by generative pre-training*.

Raworth, K. (2017). Why it's time for Doughnut Economics. *IPPR Progressive Review*, *24*(3), 216–222. https://doi.org/10.1111/newe.12058

Roberts, H., & Hine, E. (2023, September 27). *The future of AI policy in China | East Asia Forum*. https://eastasiaforum.org/2023/09/27/the-future-of-ai-policy-in-china/

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, *65*(6), 386–408. https://doi.org/10.1037/h0042519

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536. https://doi.org/10.1038/323533a0

Sharma, S. (2024, December 26). DeepSeek-V3, ultra-large open-source AI, outperforms Llama and Qwen on launch. *VentureBeat*. https://venturebeat.com/ai/deepseek-v3-ultra-large-open-source-ai-outperforms-llama-and-qwen-on-launch/

Sheehan, M. (2023, July 10). *China's AI Regulations and How They Get Made*. Carnegie Endowment for International Peace. https://carnegieendowment.org/research/2023/07/chinas-ai-regulations-and-how-they-get-made?lang=en

Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., & Anderson, R. (2023). The Curse of Recursion: Training on Generated Data Makes Models Forget. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2305.17493

Smith, B. (2025, January 3). *The Golden Opportunity for American AI*. Microsoft On the Issues. https://blogs.microsoft.com/on-the-issues/2025/01/03/the-golden-opportunity-for-american-ai/

Sokler, B., Hecht, A., & Tamotsu Fjeld, C. (2024, December 19). *FY 2025 NDAA Passes with AI Provisions, Awaiting Biden's Signatur*. https://natlawreview.com/article/congress-passes-defense-bill-ai-provisions-ai-washington-report

State of California Department of Justice. (2025, January 13). *Attorney General Bonta Issues Legal Advisories on the Application of California Law to AI*. State of California - Department of Justice - Office of the Attorney General. https://oag.ca.gov/news/press-releases/attorney-general-bonta-issues-legal-advisories-application-california-law-ai

Strickland, E. (2024, April 15). *Stanford's 2024 AI Index Tracks Generative AI and More—IEEE Spectrum*. https://spectrum.ieee.org/ai-index-2024

Terrell, M. (2024, October 14). *New nuclear clean energy agreement with Kairos Power*. Google. https://blog.google/outreach-initiatives/sustainability/google-kairos-power-nuclear-energy-agreement/

Text - H.R.5009 - 118th Congress (2023-2024): Servicemember Quality of Life Improvement and National Defense Authorization Act for Fiscal Year 2025 (2023-07-27). (2024, December 20). [Legislation]. https://www.congress.gov/bill/118th-congress/house-bill/5009/text

Text - H.R.6216 - 116th Congress (2019-2020): National Artificial Intelligence Initiative Act of 2020 (2020-03-12). (2020, March 12). [Legislation]. https://www.congress.gov/bill/116th-congress/house-bill/6216/text

Text - H.R.6395 - 116th Congress (2019-2020): William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021 (2020-03-26). (2021, January 1). [Legislation]. https://www.congress.gov/bill/116th-congress/house-bill/6395/text

The Republican Party Platform, 2024. (2024, July 15). Ballotpedia. https://ballotpedia.org/The_Republican_Party_Platform,_2024

The White House. (1995, November 8). *Technology and Economic Growth*. https://clintonwhitehouse5.archives.gov/textonly/WH/EOP/OSTP/html/techgrow.html#intro

The White House. (2023a, July 21). *FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI*. The White House. https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/

The White House. (2023b, October 30). *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*. The White House. https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

The White House. (2023c, November 1). *FACT SHEET: Vice President Harris Announces New U.S. Initiatives to Advance the Safe and Responsible Use of Artificial Intelligence*. The White House. https://www.whitehouse.gov/briefing-room/statements-releases/2023/11/01/fact-sheet-vice-president-harris-announces-new-u-s-initiatives-to-advance-the-safe-and-responsible-use-of-artificial-intelligence/

The White House. (2025a, January 13). *FACT SHEET: Ensuring U.S. Security and Economic Strength in the Age of Artificial Intelligence*. The White House. https://www.whitehouse.gov/briefing-room/statements-releases/2025/01/13/fact-sheet-ensuring-u-s-security-and-economic-strength-in-the-age-of-artificial-intelligence/

The White House. (2025b, January 14). *Executive Order on Advancing United States Leadership in Artificial Intelligence Infrastructure*. The White House.

https://www.whitehouse.gov/briefing-room/presidential-actions/2025/01/14/executive-order-on-advancing-united-states-leadership-in-artificial-intelligence-infrastructure/

The White House. (2025c, January 20). *Initial Rescissions Of Harmful Executive Orders And Actions*. The White House. https://www.whitehouse.gov/presidential-actions/2025/01/initial-rescissions-of-harmful-executive-orders-and-actions/

The White House. (2025d, 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/

Tobey, D., Carr, A., Bigg, C., Fulton, S., Ge, A., & Hoffner, K. (2024, September 12). *China-releases-AI-safety-governance-framework | DLA Piper*. https://www.dlapiper.com/en-us/insights/publications/2024/09/china-releases-ai-safety-governance-framework

Tremayne-Pendelly, A. (2024, December 4). Amid the A.I. Boom, These States Have Become Data Center Hubs. *Observer*. https://observer.com/2024/12/ai-demand-where-data-centers-using-most-energy/

Turing, A. M. (1950). I.—COMPUTING MACHINERY AND INTELLIGENCE. *Mind*, *LIX*(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433

United States Congress. Senate. (2024). *Driving U.S. innovation in artificial intelligence: A roadmap for artificial intelligence policy in the United States Senate*. United States Senate.

U.S. Chamber of Commerce Institute for Legal Reform. (2024, May 15). What Is a Private Right of Action. *ILR*. https://instituteforlegalreform.com/blog/what-is-a-private-right-of-action/

U.S. Congress. Committee on the Judiciary Subcommittee on Antitrust. (2020). *Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations, Part I. Congressional Publications*.

U.S Department of Energy. (2024, December 20). *DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers*. Energy.Gov.

https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers

U.S. Environmental Protection Agency. (n.d.). *Our History* | *ENERGY STAR*. Retrieved January 22, 2025, from https://www.energystar.gov/about/how-energy-star-works/history

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). *Attention Is All You Need* (No. arXiv:1706.03762). arXiv. http://arxiv.org/abs/1706.03762

Wang, Y., Qu, W., & Ye, X. (2024). *Selecting Between BERT and GPT for Text Classification in Political Science Research* (No. arXiv:2411.05050; Version 1). arXiv. https://doi.org/10.48550/arXiv.2411.05050

Warren S. McCulloch & Walter Pitts. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133. https://doi.org/10.1007/bf02478259

Weise, K., & Tamayo, J. (2024, December 25). An A.I. Boom Makes Electricians Flock to Central Washington. *The New York Times*. https://www.nytimes.com/2024/12/25/technology/ai-data-centers-electricians.html

Wiesinger, J., Marlow, P., & Vuskovic, V. (n.d.). *Agents*. Retrieved January 7, 2025, from https://www.kaggle.com/whitepaper-agents

Wiggers, K. (2025, January 11). Researchers open source Sky-T1, a “reasoning” AI model that can be trained for less than $450. *TechCrunch*. https://techcrunch.com/2025/01/11/researchers-open-source-sky-t1-a-reasoning-ai-model-that-can-be-trained-for-less-than-450/

World Nuclear News. (2024, December 18). *Oklo signs power agreement with data centre developer*. World Nuclear News. https://world-nuclear-news.org/articles/oklo-signs-power-agreement-with-data-centre-developer

Wyn Davies, C., & Dennis, G. (2024, December 20). *Getty Images v Stability AI: The implications for UK copyright law and licensing*. Pinsent Masons. https://www.pinsentmasons.com/out-law/analysis/getty-images-v-stability-ai-implications-copyright-law-licensing

Yee, L., Chui, M., Roberts, R., & Xu, S. (2024, July 24). *Why AI agents are the next frontier of generative AI | McKinsey*. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai

Zhao, H., Chen, Q. P., Zhang, Y. B., & Yang, G. (2024). *Advancing Single- and Multi-task Text Classification through Large Language Model Fine-tuning* (No. arXiv:2412.08587; Version 1). arXiv. https://doi.org/10.48550/arXiv.2412.08587