



Date: March 15, 2025

To: Office of Science and Technology Policy (OSTP);
Networking and Information Technology Research and Development (NITRD)
National Coordination Office (NCO)

From: Nicholas P. Garcia, Senior Policy Counsel, Public Knowledge

Re: Request for Information on the Development of an AI Action Plan

Dear Administration Officials,

Public Knowledge is grateful for the call for public comments and the opportunity to provide input on this Request for Information (RFI) on the Development of an Artificial Intelligence (AI) Action Plan. Public Knowledge is a non-profit, non-partisan, public interest organization that advocates for technology policies that promote free expression, protect consumer rights, and prevent monopolization of emerging digital markets.¹

Innovation through openness and competition.

As AI technology advances, ensuring its development remains open, transparent, and competitive is paramount. The Trump Administration has emphasized innovation as a key policy priority for “sustaining and enhancing America's AI dominance in order to promote human flourishing, economic competitiveness, and national security.”² These goals align with Public Knowledge's core mission, and our work on AI policy has always aimed at promoting AI innovation and ensuring its value is realized by the American people. Our analysis of the AI industry, economic landscape, and history of technology shows that the best path to AI innovation is through openness and competition, rather than control by a few dominant firms.

To ensure an open and competitive AI sector, we propose five key policy priorities:

1. Protect the rights to read and learn that make AI training possible.
2. Support open-source, collaborative research and development.
3. Build and maintain public physical and digital AI infrastructure to prevent monopolization and private enclosure.
4. Develop standards and sensible rules around AI explainability and transparency to ensure trust and adoption.

These policy priorities will not only secure American competitiveness in AI, but help create a truly competitive and innovative AI industry in the United States.

¹ <https://publicknowledge.org/about-us/>

² RFI, <https://www.federalregister.gov/d/2025-02305/p-11>.

Protect the Rights to Read and Learn that Make AI Training Possible

At the center of our vision for a creative and connected future for all are the rights to read, learn, and share knowledge. These rights not only underpin our shared commons of knowledge but also drive innovation by fostering an ecosystem where ideas and knowledge can circulate freely and openly. The protection of these rights against restrictive intellectual property rights expansions is crucial for maintaining the dynamism of the open research ecosystem that has historically propelled technological and creative advancements. This is particularly important for our burgeoning AI research ecosystem because of the critical role the collection and analysis of data plays in modern machine learning techniques.

Our stance is clear: the existing copyright framework, while not flawless, is sufficiently robust to handle the new dynamics introduced by generative artificial intelligence (GAI) and breakthroughs in machine learning. As we have previously written, AI training on copyrighted material should be permissible under our existing law.³ Our copyright system—and its limitations and exceptions—is designed to balance between protecting creators and fostering public access to new knowledge. Any overreach in copyright reforms would disrupt this delicate balance, hindering the innovative potential of AI technologies and also exacerbating existing economic disparities within creative industries.

The Risks of Copyright Expansion

An influential consortium of media companies, buoyed by popular concerns over the real risks of economic disruption and labor displacement created by AI, has agitated for expanded copyright protections in response to the advent of new AI technologies. However, expansion risks impairing the very mechanisms of learning and sharing that are critical, not just for generative AI development, but for all forms of computational research and analysis, as well as forms of human learning and knowledge sharing. Expanding the interpretation of copyright to restrict how we can read and learn from creative works would undoubtedly limit the essential activities of fair use and free expression which are crucial for a healthy intellectual and creative economy.

Moreover, expanding copyright does not address the root economic challenges faced by creators; rather, it will entrench the interests of powerful entities at the expense of smaller players and innovators. Requiring expensive licensing regimes for AI training will drive up the barriers to entry for AI research and development, driving

³ <https://publicknowledge.org/policy/usco-ai-and-copyright-comments/>;
<https://publicknowledge.org/generative-ai-is-disruptive-but-more-copyright-isnt-the-answer/>

small businesses and nonprofit actors out of the sector. Similarly, media companies being able to bargain over their troves of copyrighted content will only incentivize further media consolidation, exacerbating the monopsonistic quality of creative industries and further increasing the economic precarity and disempowerment of creative workers.

Thus, expanding copyright protection to cover AI training would contradict the principles of competition and innovation that should guide our approach to AI policy.

AI Training is Noninfringing

The application of our current copyright laws to the practice of AI training confirms that these activities align well with the objectives of copyright, which aims to stimulate creativity and promote the public good.

AI training involves creating massive datasets from digitized content, transforming these materials from their original purposes to serve as inputs for machine learning—a use that is fundamentally different from their original consumption and does not compete in the original marketplace. This data gathering exemplifies the principle of transformation, by repurposing copyrighted content in a way that adds value to society without displacing the original works.

Similarly, the process of training itself results in a new artifact, the AI model, which has a generalized and distinct use from any of the individual works used in its creation. While an AI model can be used to create new work similar to the works used in its training, this is the same kind of market competition that arises from anyone reading, learning from, and subsequently embarking on a similar enterprise.

Finally, the concern that AI might infringe on creative rights through mimicking or stylistic replication is not best addressed through the lens of copyright law. Copyright protection only extends to the specific expressive elements of a work. Instead, AI models—as well as human creators influenced by the work of others—leverage the underlying ideas, styles, or methods which are free for all to use. This critical limitation is core to the constitutionality of copyright law, protects free expression, and thus promotes innovation and vibrant creativity.

In sum, we encourage the administration to approach copyright law in the context of AI by prioritizing the freedoms that support innovation and competition. Protecting the rights to read, learn, and develop new technologies ensures that the AI landscape remains dynamic and inclusive. By resisting unnecessary expansions of copyright that could stifle these freedoms, we can foster an environment where AI enhances societal, creative, and economic opportunities for all. This balanced

approach not only aligns with the principles underlying our copyright system but also supports the broader objectives of promoting innovation and competitive fairness in the burgeoning field of artificial intelligence.

Support Open-Source, Collaborative Research and Development

The advancement of artificial intelligence through open-source methodologies is a strategic necessity in today's technology landscape. Open-source AI development fosters a collaborative environment that is essential for innovative breakthroughs and ensures widespread sharing of those breakthroughs to advance the state of the art. This approach aligns with Public Knowledge's long-standing advocacy for policies that promote technological openness and transparency, as detailed in our comments submitted to the National Telecommunications and Information Administration (NTIA) proceeding on Widely Available Model Weights (WAMW Comment).⁴

Advantages of Openness in AI

Open-source AI development is instrumental for driving safety, innovation, and competition within the tech industry. Open-source AI models allow for widespread scrutiny by researchers and developers across the globe, which enhances the safety of these technologies by exposing potential flaws and vulnerabilities that closed systems might conceal. This transparency not only leads to more robust AI systems but also accelerates the pace of innovation by allowing developers to build upon each other's work without the barriers imposed by proprietary technologies.

The competitive advantages of open-source AI are significant. By lowering barriers to entry, open-source AI allows startups and smaller tech companies to compete with larger entities, thereby preventing market monopolies and encouraging a diverse technological ecosystem. This diversity is crucial for spurring creative solutions to complex problems and for preventing any single entity from establishing too much control over AI technologies.

Openness on a Spectrum

Recognizing that openness exists on a spectrum is vital for understanding the varied benefits that different degrees of openness can provide. As noted in our WAMW Comment, openness should be defined for AI systems to promote free, open, and accessible technology and AI policy should recognize a broad spectrum of openness for AI systems.⁵ This nuanced view acknowledges that while full openness is ideal for

⁴ <https://publicknowledge.org/policy/final-ntia-comments/>

⁵ WAMW Comment at pg. 6-7.

maximizing the aforementioned benefits, any move towards greater transparency can yield significant advantages.

For instance, even partial openness, such as sharing model weights or training dataset information, can contribute positively to the AI community. This level of openness can help balance proprietary interests with public benefits, providing a gateway for more companies and individuals to engage with and contribute to AI advancements.

Encouraging More Openness

To foster an environment where more openness in AI is encouraged, it is crucial to support policies that facilitate access to public datasets, open APIs, and collaborative platforms. As discussed further below, investing in public infrastructure that supports open-source AI—such as cloud services, computing power, and storage—can democratize access to the necessary tools for open-source AI development. Furthermore, government and educational institutions can play a pivotal role by advocating for and adopting open-source AI solutions, setting a standard for transparency and collaboration.

Moreover, creating incentives for companies to adopt open-source practices can accelerate the shift towards a more open AI landscape. These incentives could include tax breaks, grants, or public recognition for companies that contribute to open-source projects or that make their proprietary technologies partially open.

In conclusion, supporting open-source and collaborative AI development is essential for ensuring that AI technologies are safe, innovative, and competitive. By advocating for more openness, whether complete or partial, we can cultivate an AI ecosystem that is not only technologically advanced but also equitable and accessible to all. This commitment to openness will underpin the continued growth and success of the AI field, ensuring it remains a positive force for technological progress in society.

The Government Should Tackle the Physical and Digital Public Infrastructure Obstacles to American AI Innovation and Leadership

Artificial intelligence has the potential to transform industries, drive economic growth, and enhance public services in ways that parallel some of the most consequential technological advancements in history. However, realizing AI's full promise requires a strong foundation—one that is currently lacking in the United States due to a lack of public investment in essential AI infrastructure. Just as past infrastructure projects in electrification, transportation, and telecommunications laid the groundwork for American leadership in the industrial and digital economies, the

federal government must now invest in “Public AI” infrastructure to ensure the nation remains at the forefront of innovation and competition.⁶

Without decisive action, AI development will remain concentrated in the hands of a few dominant corporations, reinforcing monopolistic control over computing power, data resources, and development expertise. This concentration of AI capabilities stifles competition, inhibits broader participation in AI research, and risks ceding strategic technological leadership to other nations that are already making significant public investments. A robust Public AI infrastructure strategy can counteract these trends by creating a competitive and innovative AI ecosystem that benefits businesses, researchers, and the public at large.

Historical Precedents for Public Infrastructure Investments

The federal government has historically played a pivotal role in fostering technological advancements through large-scale public investments. The electrification of America in the early 20th century was not merely a product of private enterprise but a result of government-supported initiatives like the Rural Electrification Act.⁷ Similarly, the construction of the interstate highway system in the 1950s—one of the largest public works projects in history—enabled the rapid expansion of the automotive industry and facilitated economic growth across the country.⁸ These investments did not “pick winners and losers” in the private sector but instead created a level playing field that allowed businesses of all sizes to thrive.

AI requires a similar commitment to public infrastructure. Just as roads, bridges, and power grids are essential for economic activity, AI development depends on a robust foundation of computing resources, data availability, and research capabilities. However, unlike previous infrastructure projects, Public AI infrastructure must account for both physical (compute and hardware) and digital (data and software) elements. Most importantly, these investments must be truly public, with robust public accessibility and public accountability mechanisms, ensuring that taxpayer dollars actually provide long-term benefits rather than merely subsidizing private-sector expansion.

The Risks of Privatized Digital Infrastructure

A cautionary tale can be found in the privatization of the internet. While the U.S. government initially supported the research and development of the internet through national laboratories and public institutions,⁹ much of the internet’s physical

⁶ See, <https://publicai.network/whitepaper/>

⁷ <https://livingnewdeal.org/glossary/rural-electrification-act-1936/>

⁸ <https://www.archives.gov/milestone-documents/national-interstate-and-defense-highways-act>

⁹ <https://www.nsf.gov/impacts/internet>

infrastructure was later privatized, leading to monopolization in broadband deployment. The result has been a market characterized by slow speeds, limited competition, and high costs for consumers.¹⁰ The failures of broadband policy—including uneven access, weak customer protections, and underinvestment in infrastructure—serve as a stark warning for AI. In a sector where there is already considerable market entrenchment in terms of compute power and data access, an AI strategy that does not strongly emphasize public infrastructure will follow an even sharper downward trajectory, leading to re-entrenchment and concentrated power in the hands of a few dominant Big Tech firms.

To avoid repeating these mistakes, the federal government must take proactive steps to ensure key AI infrastructure remains a public good, rather than an exclusive asset controlled by a handful of firms. This means making strategic investments that build public AI infrastructure while ensuring broad access to essential AI capabilities.

Three Key Areas for Public Investment

To foster a more open, competitive, and innovative AI ecosystem, the federal government should invest in three critical areas: public compute resources, open data infrastructure, and Public AI model development. These investments will provide the essential building blocks for AI research and development, ensuring that AI innovation is not limited to well-funded tech giants.

1. Public Compute Infrastructure: A Smarter Approach Than Simply Subsidizing Private Data Centers

One of the most pressing challenges in AI development is access to computational power. The cost of training frontier AI models is one of the primary bottlenecks in the industry, with the leading AI firms spending billions on proprietary data centers. Without access to similar compute resources, startups, research institutions, and smaller companies struggle to compete, leading to further industry consolidation at the whims of Big Tech hyperscalers.

Instead of offering subsidies or regulatory support to private tech firms to expand their existing infrastructure, the government should invest directly in public compute resources. This could take the form of national AI research compute centers housed within federal agencies, universities, or national laboratories or even state and local programs to develop community-based compute resources through digital anchor institutions like schools and libraries. By leveraging existing federal resources and research facilities, the government can create public alternatives to

¹⁰ See e.g., <https://www.npr.org/transcripts/899472976>

private-sector compute monopolies, ensuring that AI capabilities remain accessible to a diverse range of innovators and users.

Public compute initiatives like the National AI Research Resource (NAIRR) represent an important step in this direction for expanding access for academics and researchers and administration support for the broadly bipartisan CREATE AI Act would be an excellent starting place, but more ambitious efforts are needed. The Trump Administration should go further and develop policy pilots that expand compute capacity and access to entrepreneurs, small businesses, and individuals.

Similar to how public roads enabled private sector growth without picking corporate winners, publicly accessible compute infrastructure will provide a baseline level of access to AI resources, fostering competition and accelerating AI research across multiple sectors.

2. Open Data Resources: Strengthening National AI Research Capacity

Data is the lifeblood of AI. Yet, access to high-quality, well-curated datasets is restricted by cost and fears about copyright liability risk. This creates an uneven playing field where only the biggest (or most reckless) companies can access the data required to develop competitive AI models. This concern is particularly important given the ongoing race for AI dominance with China: data access, though the details remain opaque, has been cited as one of the primary advantages for Chinese AI developers.¹¹

To address these imbalances, the government should expand support for public data repositories, ensuring that high-quality datasets are available for academic and commercial research. Again, the NAIRR provides a strong starting point, but additional efforts could include expanding the role of the Library of Congress in curating publicly available digital datasets, developing AI-ready datasets through partnerships with universities and research institutions, and ensuring that copyright law does not restrict the ability to use publicly available information for AI training. Unnecessary expansions of copyright law that restrict access to factual information would not only hinder AI development but would also undermine the broader principles of knowledge-sharing and scientific progress.

By investing in open data infrastructure, the government can prevent AI knowledge from becoming the exclusive property of a few corporations, ensuring that publicly funded data resources remain accessible to all.

¹¹ See e.g., <https://www.brookings.edu/articles/how-to-tackle-the-data-collection-behind-chinas-ai-ambitions/>

3. Building Instead of Buying: Developing Publicly-Owned AI Models and Applications

For decades, the federal government has played a crucial role in advancing technological innovation, often laying the foundation for private-sector breakthroughs. However, in artificial intelligence, the government has largely limited its role to procurement rather than direct development, creating unnecessary dependence on private vendors and exposing critical public functions to monopolistic control. Expanding the government's capacity to develop AI models through national laboratories, universities, and federal agencies is essential to ensuring that AI development remains transparent, accountable, and aligned with national interests. Initiatives like the Department of Energy's FASST (Fair, Accountable, Secure, and Safe Technology) AI program demonstrate that publicly funded research can produce cutting-edge AI systems designed for public benefit rather than private profit.¹² Strengthening these efforts will ensure that AI remains an accessible and equitable technology rather than a privatized asset controlled by a few dominant firms.

Relying on private-sector AI procurement poses several risks. First, private AI models are driven by profit incentives, often neglecting applications that serve critical public needs such as scientific research, healthcare, and equitable access to AI-driven services. Dependence on private firms also leads to vendor lock-in, reducing government flexibility and increasing long-term costs. Once locked into proprietary systems, agencies face high barriers to switching providers or modifying models to meet evolving public needs. Additionally, private-sector AI systems often lack transparency, making it difficult for regulators and the public to scrutinize their decision-making processes, ensure fairness, and prevent discriminatory outcomes. On the other hand, developing publicly owned AI models enhances transparency and oversight, ensuring that government-deployed AI aligns with democratic values and public accountability standards. Federally developed models can incorporate rigorous data governance, algorithmic transparency, and auditing mechanisms, particularly for sensitive applications such as law enforcement or public health.

Public investment in AI model development can also foster competition by providing alternatives to dominant private models, reducing barriers to entry for startups, academic researchers, and smaller firms that lack access to commercial models. Government-developed models can be open-sourced or otherwise be made publicly available to ensure that there are low-cost, public options for AI models in order to create stable alternatives to private offerings.

¹² <https://publicknowledge.org/policy/comments-in-response-to-doe-request-for-information-on-fasst/>

Beyond competition and transparency, public investment in government AI research strengthens the federal workforce and enhances government expertise in AI governance. Expanding AI research and development within national labs and federal agencies cultivates a skilled public-sector workforce capable of managing, deploying, and regulating AI technologies effectively. Programs like FASST provide a framework for building internal AI capacity, reducing reliance on private contractors and ensuring that AI policies and models remain under democratic oversight.

Public AI initiatives should span the entire AI technology stack, from compute resources to datasets and pretrained models, ensuring that critical bottlenecks are identified and addressed. By investing in AI model development rather than simply purchasing proprietary systems, the government can prevent monopolization, reduce dependency on opaque private models, and build the expertise in AI technology that our government and workforce needs.

Explainability and transparency standards are essential for AI innovation and adoption.

The widespread adoption of artificial intelligence requires public confidence in its fairness, reliability, and accountability. Without clear mechanisms to understand, audit, and assess AI models, both innovation and public trust in AI technologies will be severely limited. Standards for AI transparency and explainability—established through voluntary, consensus-driven processes and coordinated by the federal government through agencies like the National Institute of Standards and Technology (NIST) and the U.S. Artificial Intelligence Safety Institute (AISl)—are essential for ensuring that AI systems are both safe and effective. By setting clear expectations for transparency and explainability, these standards can serve as the foundation for trust in AI-driven products and services, ultimately accelerating AI adoption across industries.

Transparency as a Prerequisite for AI Accountability

Transparency in AI should go beyond basic disclosures and include structured accountability mechanisms that provide meaningful insights into how AI models function, what data they rely on, and how they impact users. Public Knowledge has consistently emphasized that AI systems must not operate as black boxes but should be subject to rigorous transparency and explainability requirements to ensure accountability and fair outcomes.

AI transparency is a key component of consumer protection and business confidence. When AI models influence decisions in hiring, lending, healthcare, and

other critical domains, affected individuals and businesses should have access to clear, understandable explanations of how those decisions are made. Public disclosure of AI decision-making criteria, model training data sources, and algorithmic biases fosters trust and prevents discriminatory or harmful outcomes. Moreover, transparency should, at a minimum, be a strongly incentivized industry norm, reinforced through structured reporting standards, third-party audits, and disclosure mechanisms.

By establishing voluntary and consensus-driven transparency standards through NIST and AISI, the federal government can help businesses and developers adopt best practices that increase consumer confidence without stifling innovation. Transparency empowers users, regulators, and businesses alike by providing the information necessary to assess risks, identify biases, and ensure AI systems are functioning as intended.

Explainability as a Foundation for AI Safety and Innovation

Similarly, explainability is critical for both users and developers to understand how AI systems operate and why they make certain decisions. A well-calibrated explainability framework ensures that AI is not only interpretable to experts and regulators but also comprehensible to end users who interact with AI-powered tools. Without explainability, AI models risk being seen as opaque, unaccountable, and untrustworthy, limiting their adoption in key sectors such as finance, healthcare, and legal decision-making.

For businesses, explainability serves as a competitive advantage—companies that can offer clear, understandable AI systems will be better positioned to attract customers and comply with evolving regulatory expectations. For developers, explainability improves the ability to debug, refine, and enhance AI models, leading to more reliable and adaptable AI systems over time. Yet proof of these capabilities remains technically challenging. Further developing research into AI explainability and developing an accountability environment that can validate AI behavior should be a priority of this administration. Public Knowledge has emphasized that AI accountability mechanisms—including certifications, audits, assessments, and testing—should be more than just “trust seals”. Instead, they must be meaningful, enforceable, and integrated into broader safety and consumer protection frameworks.¹³

To ensure effective AI adoption, both explainability and transparency standards should be developed in coordination with industry leaders, researchers, and consumer advocacy groups, ensuring that they are both technically feasible and

¹³ <https://publicknowledge.org/policy/ntia-ai-accountability/> at pg. 3.

aligned with public interests. Agencies like NIST and AISI should play a leading role in establishing best practices, testing explainability methodologies, and guiding businesses toward responsible AI deployment. AI innovation and adoption will only reach their full potential if businesses, consumers, and regulators trust that AI systems operate fairly, make accountable decisions, and can be independently assessed. Voluntary and consensus-driven standards-setting efforts led by NIST and AISI will provide the framework needed to ensure transparency and explainability while maintaining flexibility for industry advancement. These standards should encourage widespread disclosure, independent testing, and structured auditing mechanisms that enhance both AI safety and public confidence. By prioritizing transparency and explainability, the U.S. can position itself as a leader in responsible AI governance, ensuring that AI-driven technologies remain trustworthy, competitive, and aligned with democratic values.

Conclusion

Artificial intelligence presents extraordinary opportunities for economic growth, scientific discovery, and public benefit, but its continued success depends on maintaining an ecosystem that is open, competitive, and accountable. To ensure that AI innovation is aligned with public interest, this administration should take proactive steps to protect the rights to read and learn, support open-source AI development, invest in public infrastructure, and develop transparency and explainability standards that build trust in AI systems.

To ensure innovation and American AI leadership on the global stage, Public Knowledge urges the administration to prioritize policies that foster competition, encourage collaboration, and prevent industry consolidation from stifling the potential of this potentially revolutionary technology.

By embracing competition and openness, and taking a bold lead with Public AI infrastructure investment, the U.S. will ensure it continues to lead the world in AI innovation. Public Knowledge looks forward to opportunities to advise the administration on developing an AI Action Plan and policies that ensure AI benefits all Americans.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.