

# PUBLIC SUBMISSION

<b>Received:</b> May 29, 2025 <b>Tracking No.</b> mba-7poz-rp8f <b>Comments Due:</b> May 28, 2025 <b>Submission Type:</b> API
--

**Docket:** NSF-2025-OGC-0001  
NITRD\_FRDOC\_0001

**Comment On:** NSF-2025-OGC-0001-0001  
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

**Document:** NSF-2025-OGC-0001-DRAFT-0339  
Comment on FR Doc # 2025-07332

---

## Submitter Information

**Organization:** Adobe

---

## General Comment

See attached file(s)

---

## Attachments

Adobe Comments -OSTP AI Research and Development RFI

Adobe Comments - OSTP-NSF-AI Research and Development RFI

Adobe Comments - NSF-AI Research and Development RFI

# **Adobe Comments on the National Science Foundation's Request for Information on the Development of a National Artificial Intelligence Research & Development Strategic Plan**

## **Introduction**

Adobe welcomes the Administration's efforts to establish a forward-looking U.S. policy for sustaining and enhancing artificial intelligence ("AI") innovation in the U.S. in order to promote human flourishing, economic competitiveness, and national security. We are living in the midst of an incredible era of technological transformation. AI represents a powerful force that will lead to unprecedented levels of productivity, creativity, and economic growth across the country. By embracing this moment, the U.S. can lead the world in shaping an AI-powered future that drives the next era of innovation and economic prosperity.

We appreciate the opportunity to submit recommendations to the National Science Foundation (NSF) and the Office of Science and Technology Policy (OSTP) via the Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO) on how government can play a vital role in advancing AI research and development. By continuing to invest in this cutting-edge technology and establishing robust industry standards, we can create an environment that will unlock AI's full potential and build consumer confidence—paving the way for broad adoption and long-term success.

One essential opportunity ripe for government leadership is the development of a first-of-its kind AI testing pilot and standardized benchmarking program which would enable companies to evaluate their models through a scalable, industry-driven framework that fosters public trust, enables innovation, and drives consumer adoption of AI. Importantly, standardized testing and benchmarking would also ensure that AI systems work equitably, free from ideological bias or engineered social agendas. As we detail below, Adobe believes this is an important effort the government can lead by using its power to convene critical perspectives; promote competition in the AI ecosystem; and enable consumer choice and innovation, ensuring continued U.S. leadership on the global AI stage.

## **Overview of Adobe**

At Adobe, our mission is to change the world through personalized digital experiences. Since our founding in December 1982, we have pioneered transformative technologies that allow our customers—who range from emerging artists to global brands—to channel their imaginations, unleash their creativity, and power their businesses.

As technology evolves and becomes more advanced, we are committed to innovating responsibly and advancing the responsible use of technology. This means making sure our technologies and the processes we go through to develop them are accountable, responsible,

and transparent. It also means considering the broader implications of new technologies and stepping up as leaders on societal issues where Adobe can leverage our technology and expertise to have a unique impact.

Adobe's business is comprised of a comprehensive technology suite that fuels both creativity and productivity. Across our product portfolio, Adobe is building on a decade-long legacy of AI innovation by leveraging the power of AI to transform ideation, creation, and production. For business professionals and consumers, digital document products like Adobe Acrobat and the all-in-one visual creation tool Adobe Express power end-to-end solutions for creating, sharing and consuming content. For professional creatives and creators, tools such as Illustrator and Photoshop provide the power and precision needed to bring creative visions to life across any media type and surface. And for marketers who need to deliver personalized digital experiences at scale, Adobe Experience Platform helps to reach the right audiences and drive meaningful business outcomes.

In March 2023, we launched [Adobe Firefly](#), our family of creative generative AI (GAI) models that are commercially safe and developed with respect for creators' rights. Firefly models allow users to channel their creativity in entirely new ways and are available both in the standalone Firefly app and seamlessly integrated with Adobe applications such as Photoshop, Lightroom, Illustrator, Premiere Pro and Adobe Express. By unifying image, video, audio and vector generation and providing unmatched creative control, Firefly models allow creators to work more productively and with an unmatched degree of precision. We continue to enhance our Firefly models by improving quality and performance, adding the most requested features from the community, and incorporating them when it can complement and improve workflows.

We believe that AI done right will amplify human creativity and capabilities to new levels with deeper insights, accelerated task performance, and improved decision-making ability. As we continue to harness the power of AI, we are committed to developing and deploying AI in line with our [AI Ethics principles](#) of Accountability, Responsibility, and Transparency, while considering its broader impact on society.

### **Recommendation: Invest in the Development of an AI Testing Pilot and Standardized Benchmarking Program**

#### **Training and Testing Data is Essential to AI Innovation**

An essential component of AI development is the data on which it is trained. The more information or data you give an AI model, the better it will perform. This is why AI developers have put so much emphasis on enabling access to high-quality data in order to create high-performing, accurate, safe and secure AI models.

The ideal outcome is one where AI developers have access to enough high-quality data to effectively develop world-leading AI models without collecting more than what is necessary. One way AI developers can continue to augment their models and ensure they are ready for

deployment is by continuously testing them using benchmark test datasets and well-defined metrics representing acceptable ranges of outcomes.

Today, however, there is no consistency across industry regarding what is considered “acceptable” or what areas to test against. In the absence of clearly established standards that developers can rely on, the industry is at risk of instead becoming subject to burdensome requirements and government audits, which may not even accurately mitigate harms they seek to mitigate, but which will almost certainly decelerate AI innovation.

### **Leading the World with an AI Testing Pilot and Standardized Benchmarks**

Rather than imposing burdensome government requirements that could restrict private sector AI development and deployment, the administration has an opportunity to lead an alternative approach, unique to the U.S., by establishing an AI Testing Pilot Program to ensure that users, including younger users, can interact safely and appropriately with AI systems. The AI Testing Pilot would include convening various stakeholders including members of industry to develop robust benchmark test datasets and standards that companies can use to evaluate their AI models.

Industry-led standards provide uniform guidelines to ensure quality, safety, and interoperability for AI systems, ensuring consistency across different verticals and industries. At the same time, standards encourage responsible development, which allows companies to build trust with their customers, by giving companies clear guidelines without imposing unnecessary restrictions. Finally, a set of standards will help achieve the administration’s goal of developing AI systems that are free from ideological bias and engineered social agendas as stated in [Executive Order 14179](#), by collectively agreeing upon what to test for and determining what the acceptable ranges of outcomes are, rather than leaving this up to companies, individuals, or a patchwork of governments in different jurisdictions.

A collaborative effort between the public and private sectors will enable the development of benchmarks focused on allowing companies to push the boundaries of AI technology without unnecessary restrictions, while ensuring that users, including parents, have the right tools and information to make informed decisions about which AI systems to interact with or let their children interact with. By fostering an environment of transparency and accountability, this approach will build public trust in AI, ultimately leading to broad consumer adoption of this powerful technology.

### **Example Use Case: AI Image Generator**

Take an AI image generation model, for example. Here, categories to test against could include explicit content, suicide content, self-harm content, eating disorder content, and violent content. While Adobe and other companies already spend significant resources to moderate AI image output to eliminate such things as self-harm or explicit content to protect children, an independent measure of AI image output quality can provide parents, for example, with objective tools to measure safety. The AI model developer could leverage a set of testing data,

including a series of agreed upon prompts, provided by OSTP. The developer could input those prompts into the model and measure whether and how often the prompts elicited outputs that fell within those established categories of explicit content, suicide content, self-harm content, eating disorder content, hate content, and violent content.

Then, they could compare their models' results to a set of acceptable ranges agreed upon by industry for AI image generation models. While of course, AI developers should seek to minimize harms, having a range of what is acceptable reflects the reality that a stringent zero tolerance is impractical with an evolving technology and could cause developers to unintentionally block content that is actually safe. With these ranges established, if there are any areas where the AI model's results fall outside of the acceptable range, the AI developer can focus their efforts to improve their model by conducting further training and algorithmic improvements specifically for that category of output prior to deployment. In the illustrative example below, the AI image generator falls outside of the acceptable range for self-harm. This would indicate they should conduct further training and implement additional mitigations to meet the established standards.

<b>Sample AI Image Generator Internal Evaluation Report</b>		
<b>Risk category</b>	<b>Error rate</b>	<b>Acceptable range</b>
Explicit content	0.1%	< 1%
Suicide content	0%	< 1%
Self-harm content	2%	< 1%
Eating disorder content	0.25%	< 3%
Violent content	0.1%	< 3%

*\*\*Disclaimer: the above sample evaluation report was created for illustrative purposes only and is not consistent with Adobe's existing moderation programs for its AI image generation tools.\*\**

Once these categories and metrics are well-established, industry can then leverage them to help develop age-appropriate AI models, such as by using them to inform globally-accepted industry-wide ratings systems. This could be similar to how the Entertainment Software Ratings Board – an industry-led self-regulatory organization – assigns standard ratings to video games and apps so that parents can make informed decisions about which games are right for their children. By supporting this approach and encouraging industry to adopt it, the Administration can help ensure users and stakeholders make informed decisions about content and technology and whether it is suitable for specific applications or audiences, including children.

### **Advancing an Innovation-Forward Approach to AI Governance**

An industry-led standards and testing approach preserves innovation and freedom, allowing new ideas to flourish while maintaining public trust. We urge the Administration to advance this world-leading approach by leveraging its power to:

- **Convene Stakeholders across the Public and Private Sectors to Create Testing Data and Develop Benchmarks.** Industry input is essential to building practical and useful standards, classification systems, and benchmark test datasets. By collecting feedback, developing working groups, and focusing industry effort on a unified program, OSTP can facilitate the consensus-building among members of industry that an effective standard and testing program needs. This would, in turn, foster an environment for a scalable and industry-driven approach to AI governance without hampering technological progress. Rather than burdening AI companies with restrictive compliance measures, this approach would incentivize them to refine their models through transparent, standards-driven evaluation, ensuring AI development remains innovative, effective, and globally competitive. The private sector could work to establish testing data and benchmarks across various verticals, industries, markets, and use cases.
- **Promote Competition in the AI Ecosystem.** In order to support ongoing AI innovation, the U.S. must continue to create opportunities that enable AI developers of all sizes to focus their investments on cutting-edge technological development, not costly compliance with burdensome regulations. Furthermore, sourcing test data and developing testing programs and benchmarks requires significant resources. A system where each company must take this on themselves could be cost-prohibitive to emerging AI startups and narrow the AI field to just a few large companies. Instead, by maintaining the centralized repository of industry-established testing data and making it available to all AI developers, the administration can facilitate a system where companies including startups can focus their resources on developing new technology and bringing the next world-changing ideas to market. This system will boost competition in the AI marketplace, thereby spurring more innovation and continued flourishing of the AI economy.

- **Enable Consumer Choice and Innovation through a Self-Certification Process.** Once testing data has been curated and benchmarks have been established, industry can measure their AI models against these benchmarks and self-certify that they meet certain standards. Consumers can then rely on these certifications to assess different AI models leveraging an objective set of metrics, and decide which models are best suited for their needs. Government, too, can leverage them in their own procurement of AI technologies. Developers are accountable for their development of AI and have an incentive to participate, as they could create a safe harbor against future litigation. Developers could even use third-party auditors to verify that they meet the benchmarks if they choose to, creating legal certainty for companies and assuring governments and consumers that AI was soundly developed. There is a false dichotomy that exists today that stipulates that in order to develop AI responsibly, we will undoubtedly slow down innovation. But this does not have to be the case. By promoting widespread adoption of industry-led standards and the self-certification process, the Administration can in parallel empower and protect consumers by promoting the development of robust, resilient, and secure AI models and drive AI innovation. The U.S. has the opportunity to create the foundation for effective and practical AI governance globally.

## **Conclusion**

The last several years have seen an explosion in U.S. AI innovation, but this is just scratching the surface of what this industry has the power to become. By advancing a world-leading AI Testing Pilot and Standardized Benchmarks program, the Administration can continue to drive AI innovation, support the burgeoning AI startup economy, and empower consumers with more choice and control over how they interact with AI models. OSTP is uniquely positioned to leverage its decades of science, engineering, and technology policy expertise to convene industry stakeholders and facilitate the development of this important program.

Adobe appreciates the opportunity to provide input on the Administration's National Artificial Intelligence Research and Development Strategic Plan. We appreciate that the NITRD NCO has already taken important steps to consult with interested stakeholders on behalf of OSTP and the NSF as the Administration implements its AI policy priorities. We look forward to working together with the NSF, NITRD NCO, OSTP, and the White House to ensure the U.S. can achieve the Administration's goal of leading the world on AI.