Before the

National Science Foundation

Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

Comments of the Electronic Frontier Foundation

March 13, 2025

Submitted by:

Corynne McSherry
Electronic Frontier Foundation
815 Eddy Street
San Francisco, CA 94109
Telephone: (415) 436-9333 x 122
corynne@eff.org

The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF submits these comments to assist the National Science Foundation's development of an Artificial Intelligence Action Plan. AI is riding a wave of hype into adoption in a wide variety of industries and government operations.[1] While current machine learning technologies have some positive applications, they are also being adopted in consequential decision-making contexts where these emerging technologies are likely to cause harm and unlikely to deliver the promised benefits. Any action plan should prioritize identifying and mitigating such harms, while ensuring the government regulations do not unduly hamper other forms of AI innovation.

[1] 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.

A. AI, Algorithmic Decision-Making, and Transparency in AI Use and Development

The use of algorithmic decision-making tools (ADMs) by government agencies in adjudicating people's rights and privileges is of particular concern. Governments increasingly rely on algorithmic systems to make consequential assessments and determinations about people's lives, from judging eligibility for social assistance to automated and so-called "AI-enhanced" surveillance at the U.S.- Mexico border.

AI tools have been shown to be deficient when used in these sorts of complex contexts. At best, this technology can reproduce the patterns present in a training data set. At worst, it can—and often does—fail in troubling and unpredictable ways. When used to inform decisions that implicate the rights of Americans, AI reproduces historic bias by design and presents a high risk of causing new harm. Human rights violations cannot be justified by promises of mere cost savings—promises which are failing to manifest in the private sector, as workers find themselves putting in *more* labor to correct inaccuracies created by machine learning systems.

There are huge risks to using machine learning technology for criminal investigation or punishment or to determine eligibility for housing, medical care, employment, or other essential human needs. Government and private use of these systems must be regulated carefully to avoid infringing the civil rights of persons subject to their decisions. For example, we have seen media reports that the Department of Government Efficiency intends to use AI to evaluate federal workers, and use the results to make decisions about their continued employment. Such use of AI to make important decisions about people is likely to result in irrational and discriminatory employment decisions.

At the same time, government AI procurement has moved with remarkable speed. This has led to an alarming lack of transparency in government use of AI that has entrenched the largest AI companies. Without a transparent process, there is a much greater risk of wasteful spending as federal resources are poured into systems with no proven track record.

Two practices can help mitigate this risk.

The first is implementing a robust public notice-and-comment practice consistent with the Administrative Procedure Act, which requires public notice and comment for many types of agency action. Just as an agency would have to give notice and invite comment in order to change rules for deciding eligibility or action, it should be required to do so when adopting an AI or ADM tool that informs such a decision. A public and transparent notice-and-comment process will help reduce harm to the public and government waste by working to weed out bogus products and identify applications where certain types of tools, such as AI, are inappropriate.

The second is favoring technologies developed in accordance with the widely-held transparency principles of free and open-source software. By using technology that is

developed transparently and subject to adversarial review, we can ensure that the supposedly scientific basis of many ADM tools holds up to scrutiny. Abiding by core transparency principles will also enable agencies and the public to have more informed conversations about the merits and drawbacks of particular AI systems. Transparency is key because state legislatures around the country, as well as Congress, have begun to grapple with questions of fairness and legal compliance when secret AI and ADM systems are used.

It's important to note that although there is a clear need to regulate AI, policymakers should not rush to adopt a regulatory framework that would consolidate the industry by locking out small innovators. Regulating general-purpose tools too aggressively would both punish innocent actors and favor the large, incumbent companies that can afford legal battles, while pushing out academic and startup innovators. Focusing on speculative, long-term, catastrophic outcomes from AI pulls attention away from the AI-enabled harms that are directly before us. Accordingly, while those who misuse AI tools should be subject to appropriate legal constraints, any transparency framework should not unduly burden the ability of technologists, particularly small innovators, to develop general purpose AI tools just as they develop other general-purpose tools that may be used for both malicious and beneficial purposes. Regulators should focus on *the use* in question, not the tool itself.

Recommendations

1. The Action Plan should support transparency efforts in AI procurement, development and use whenever possible.

B. Copyright Concerns in Generative AI Regulation

Anxiety about generative AI is growing almost as fast as the use of the technology itself. Artists are increasingly concerned about the harms of AI tools used to mimic their respective styles. In addition to the now-infamous AI-generated song that seemed to feature Drake and The Weeknd, digital artists, musicians, actors, writers, and others are seeing their names regularly invoked, without their permission, to generate new works.

Despite the flurry of lawsuits, most new works that are created using generative AI, and the training of the tool itself probably do not infringe the copyright in any work used to train that AI tool.

That said, there are legitimate concerns that may require some rules of the road. As they consider drafting such rules, policymakers should answer some crucial questions:

- **Is the proposed regulation properly focused?** Generative AI is a category of general-purpose tools with many valuable uses; legislators should avoid technology mandates that might inhibit the development of those tools, particularly by smaller innovators that seek to compete with entrenched oligopolies.

- **Are the harms the proposal aims to alleviate documented or still speculative?** Thoughtful researchers and civil society groups have been sounding the alarm about the risks of AI-based decision-making for years. We should not let hyperbole and headlines about the *future* of generative AI distract us from addressing the damage being done by other forms of AI *today*.
- **Is the proposed regulation flexible enough to adapt to a rapidly evolving technology?** Technology often changes much faster than the law, and those changes can be difficult to predict, let alone accurately legislate around.
- **Will the rule alleviate the harm it targets?** This question gets overlooked far too often. For example, there have been several proposals to require generative AI users and developers to "watermark" the works they produce. Watermarking of AI generated content is an easy-sounding fix, but research into adversarial watermarking for AI is just beginning, and there's no strong evidence to show that it will fix the thorny problem of disinformation.
- **Finally, how does it affect other public interests?** For example, proposals designed to ensure remuneration for creators, such as a new copyright licensing regime, could make socially valuable research based on machine learning and data mining prohibitively complicated and expensive. Please see the following section for a fuller discussion of this issue. EFF has great sympathy for creators who struggle to be appropriately compensated for their work. But we must look for ways to ensure fair pay that don't limit the potential for all of humanity to benefit from valuable secondary uses.

Recommendations

1. The Action Plan should avoid embracing overly broad regulations, such as those proposed in bills such as NO FAKES and NO AI Fraud, that do not offer satisfactory answers to the questions above.

C. AI Licensing

Some have suggested that licensing schemes provide a way to address creators concenrs about use of their works for training. But any such scheme carries significant risks. Requiring developers to license the materials needed to create AI technology threatens the development of more innovative and inclusive AI models, as well as important uses of AI as a tool for expression and scientific research. Specifically, requiring AI developers to get authorization from rights holders before training models on copyrighted works would make it harder for newer companies that don't have their own trove of training data to create new tools. Instead, this scheme benefits giant tech monopolists that can afford to pay pricey licensing deals that lock in their dominant positions in the generative AI market by creating prohibitive barriers to entry.

Further, mandatory licensing through copyright is unlikely to provide any meaningful economic support for vulnerable artists and creators. Notwithstanding the highly publicized demands of musicians, authors, actors, and other creative professionals,

imposing a licensing requirement is unlikely to protect the jobs or incomes of the underpaid working artists that media and entertainment behemoths have exploited for decades. Because of the imbalance in bargaining power between creators and publishing gatekeepers, trying to help creators by giving them new rights under copyright law is, as EFF Special Advisor Cory Doctorow has written, like trying to help a bullied kid by giving them more lunch money for the bully to take.

Entertainment companies' historical practices bear out this concern. For example, in the late-2000's to mid-2010's, music publishers and recording companies struck multimillion-dollar direct licensing deals with music streaming companies and video sharing platforms. Google reportedly paid more than $400 million to a single music label, and Spotify gave the major record labels a combined 18 percent ownership interest in its now-$100 billion company. Yet music labels and publishers frequently fail to share these payments with artists, and artists rarely benefit from these equity arrangements. There is no reason to believe that the same companies will treat their artists more fairly in the AI context.