

COMMENT FOR THE RECORD

NetChoice Comments on the Jul-2024 NIST AI Safety Paper

NetChoice is a trade association of leading internet businesses that promotes the value, convenience, and choice that internet business models provide American consumers. Our mission is to make the internet safe for free enterprise and for free expression. We also work to promote the integrity and availability of the internet on a global stage, and are engaged on issues in the states, in Washington, D.C., and in international internet governance organizations.

The Transformative Potential of AI

2023 was a watershed year for artificial intelligence (AI) advancement, and 2024 is set to be the year when many of AI's promises become reality. AI has the potential to profoundly enhance our lives across a wide range of domains, from healthcare and education to productivity and creativity. However, as with any transformative technology, AI also introduces certain risks that we must anticipate and navigate.

The NIST Draft acknowledges there are many significant “challenges” to identify and measure risks when people begin to use an evolving technology like AI (pages 2-3). Key among those challenges is that “It is difficult to accurately emulate threat actors and misuse.”

However, NIST sets these challenges aside and starts with “Objective 1: Anticipate potential misuse risk”. NIST sets forth practices, recommendations, and safeguards, including “Stop development if a model displays significant misuse risk.” (p.19).

If NIST's risk-avoidance approach had been applied to the ways that Americans might have used earlier technological innovations, much of our nation's industrial base and economic

development would never have been realized. In our comment, we describe existing regulatory structures that address risks of misuse, raise concerns that premature and imprecise regulation will significantly undermine the benefits of AI, and explain why Congress alone – not NIST – has the authority to enact laws regulating development of a new technology like AI.

AI is Already Heavily Regulated

While some have called for extensive new regulations on AI, including the proposals in this draft, the reality is that this technology is already subject to a wide array of existing laws and regulatory frameworks. Any AI system must comply with the same rules as any other technology or business practice in its sector. This means that AI applications in healthcare are regulated by HIPAA and FDA guidelines, AI in finance is subject to FCRA and ECOA, and AI in education must adhere to FERPA, just to name a few examples. The notion that AI will inhabit some kind of lawless “wild west” is simply false.

For example, the federal government has already declared it a crime to intentionally lie about the time, manner, or place of an election in order to prevent qualified voters from voting. This means the government is already empowered to prosecute individuals publishing deepfakes that subvert election integrity. Moreover, existing consumer protection laws, such as the FTC Act’s prohibition on unfair and deceptive practices, already provide robust safeguards against AI systems that might mislead consumers or otherwise cause them harm.

The FTC has made clear that it will vigorously police the AI industry under its existing authorities, and has already brought enforcement actions against companies for making misleading claims about their AI products or failing to secure sensitive data used in AI development. At the same time, broadly applicable anti-discrimination statutes like the Civil Rights Act, Fair Housing Act and Americans with Disabilities Act all constrain the use of AI in high-stakes domains like employment, credit and housing to prevent disparate impacts. Finally, existing defamation and false light torts will protect the subjects of deepfake media from reputational harm.

To be clear, this is not to say that every conceivable AI harm is perfectly addressed by current law, or that thoughtful, targeted updates may not be warranted in certain areas. But the core frameworks for regulating the responsible development and use of AI are very much in place today. Policymakers and the public can take comfort in the fact that our existing legal structures are, by and large, well-equipped to prevent and remedy the highest-risk AI failures.

Before rushing to promulgate sweeping new AI-specific regulations, we should think carefully about how these new rules would interact with the dense, overlapping web of existing rules. The goal should be to strategically fill discrete gaps – not to create a redundant layer of AI law and regulation that could impede innovation while adding little marginal protection for the public.

The Dangers of Imprecise Legislative Definitions for AI

As regulators grapple with the rapid advancements in AI and its potential impact on society, it is crucial to approach the regulation of this technology with care and precision. Imprecise definitions and overly broad language in AI-related legislation can lead to unintended consequences, stifling innovation and infringing upon free speech rights.

For example, the FCC’s proposed definition of AI-generated content as “an image, audio, or video that has been generated using computational technology or other machine-based system that depicts an individual’s appearance, speech, or conduct, or an event, circumstance, or situation” suffers from similar overbreadth issues as we’ve seen in state legislation. This could require disclosures for routine video editing or audio processing that poses no real risk of voter deception.

Under the FCC’s broad definition, even basic photo editing tools, such as auto color correction or cropping, could potentially be classified as AI-generated content. As a result, a political ad featuring a candidate standing in front of a green screen or a picture that has been automatically cropped to fit a specific aspect ratio could be subject to the disclosure requirement. This would lead to a proliferation of disclaimers which would render an otherwise truthful political ad labeled with a self-declaration of “fake.” This will cause confusion among voters and undermine the effectiveness of political communication.

Moreover, the proposed penalties for non-compliance are severe. This means that broadcasters, cable operators, and other regulated entities could potentially face significant fines for minor modifications to images, even if these changes have no material impact on the message being conveyed. Such disproportionate consequences could have a chilling effect on political speech, as campaigns and advocates may be hesitant to use even the most basic digital tools for fear of running afoul of the law.

To avoid unintended consequences, the FCC and NIST must work closely with AI experts, industry stakeholders, and civil society groups to craft precise, technology-neutral definitions that focus on specific behaviors and outcomes rather than broad categories of tools. This approach will ensure that the law can adapt to the rapidly evolving AI landscape while still protecting the rights and interests of individuals and society as a whole.

The regulation of AI is a complex and delicate task that requires careful consideration and precise language. Imprecise definitions can lead to a host of unintended consequences, from chilling political speech to stifling innovation and creativity. As we navigate the challenges and opportunities presented by AI, it is essential that our laws strike a balance between protecting the public interest and fostering the responsible development and deployment of this transformative technology.

NIST Lacks Legal Authority for This Quasi-Regulatory Framework

NIST's statutory mandate under the National Institute of Standards and Technology Act (15 U.S.C. § 271 et seq.) is to develop voluntary technical standards – not to impose broad regulatory frameworks on private industry. This NIST proposal on AI safety, with its extensive prescriptive guidelines, stretches far beyond NIST's legal remit.

Quasi-Regulatory Nature of the Proposal. While NIST characterizes these guidelines as voluntary, their practical effect would be quasi-regulatory:

- **Federal Procurement:** Federal agencies often require compliance with NIST standards for procurement contracts. This effectively makes the standards mandatory for any company wishing to do business with the government in AI-related fields.

- **State and Local Adoption:** Many state and local governments adopt NIST standards as requirements, either through legislation or regulatory action. This creates a patchwork of mandatory compliance across jurisdictions.
- **Industry Standard of Care:** Courts often look to industry standards, including NIST guidelines, when determining the standard of care in liability cases. This creates de facto legal obligations for companies to follow these guidelines or risk significant legal exposure.
- **International Influence:** NIST standards often influence international standards and regulations, potentially leading to mandatory compliance requirements in other countries for U.S. companies operating globally.
- **Investor and Insurance Pressure:** Investors and insurers increasingly require compliance with NIST standards as a condition of investment or coverage, making them effectively mandatory for many businesses.

Given these government-created realities, the proposed guidelines would serve as a quasi-regulatory framework despite their nominally voluntary status. This raises serious legal concerns about NIST’s authority to issue such far-reaching directives.

Lack of Clear Congressional Authorization. As the Supreme Court recently affirmed in *West Virginia v. EPA*, 597 U.S. 697 (2022), *Biden v. Nebraska*, 600 U.S. 477 (2023), and *Loper Bright Enterprises v. Raimondo*, 603 U.S. ___ (2024), agencies cannot claim sweeping new powers without clear congressional authorization. The Court emphasized that it expects Congress to speak clearly when authorizing an agency to exercise powers of vast “economic and political significance.” NIST’s attempt to establish a de facto regulatory regime for AI development runs afoul of this principle.

The National Institute of Standards and Technology Act does not provide clear authorization for NIST to issue guidelines of this scope and impact. The Act’s primary focus is on the development of measurement standards and the promotion of industrial competitiveness. Nowhere does it contemplate NIST taking on a role akin to a regulatory agency for emerging technologies.

Federal Restrictions on AI Should Come from Congress, Not NIST

While we appreciate NIST's intent to address potential harms from AI-generated content, we believe that such far-reaching regulations should be enacted by Congress, not regulatory agencies. The proposed rules in this draft raise significant questions about the scope of NIST's authority and the appropriate balance between transparency and free speech in political discourse.

This unilateral action by a regulatory agency sets a dangerous precedent that could undermine the separation of powers and checks and balances enshrined in our Constitution. It is the responsibility of Congress to carefully consider and debate the complex issues surrounding AI regulation and to craft well-balanced legislation that addresses both the challenges and opportunities presented by this transformative technology.

Moreover, the restrictive regulatory approach outlined in the draft threatens to stifle innovation and hinder the competitiveness of the American AI industry. The introduction of burdensome and complex regulations, without proper congressional oversight, will likely discourage investment in AI research and development, as companies face increased uncertainty and compliance costs. This could lead to a slowdown in innovation, allowing other nations, such as China, to surpass the United States in the global race for AI supremacy.

Laws are Made by the Legislative Branch, Not the Executive

President Biden's recent executive order on AI¹ is not only a violation of the Constitution but also a misguided attempt to regulate a rapidly evolving technology that holds vast potential to improve people's lives. By attempting to govern AI development through an executive order, the President is effectively usurping the role of the legislature, which is the only branch of government constitutionally authorized to create such rules and regulations.

¹ White House, Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct. 30, 2023).

Moreover, the restrictive regulatory approach outlined in the executive order threatens to stifle innovation and hinder the competitiveness of the American AI industry. The executive order's broad regulatory measures will result in stifling new companies and competitors from entering the marketplace, effectively consolidating power in the hands of a few large tech giants. This not only limits consumer choice but also significantly expands the power of the federal government over American innovation.

The order puts any investment in AI at risk of being shut down at the whims of government bureaucrats, which is a dangerous approach for our global standing as the leading technological innovator.

There are already many regulations in place that govern the use of AI. Instead of examining how these existing rules can be applied to address modern challenges, President Biden chose to unilaterally expand the powers of the federal bureaucracy. This approach is not only unnecessary but also counterproductive to the responsible development of AI technology.

Only the legislative branch can craft legislation addressing the challenges and opportunities presented by AI. It is crucial that we do not allow fears to restrain the United States from improving people's lives. It is the responsibility of the legislature – not the executive branch – to ensure that any regulatory framework strikes the right balance between promoting innovation and protecting the public interest.

Conclusion

We believe the key to addressing risks and unlocking AI's full potential is to pursue a balanced, multi-stakeholder approach. We should strategically update existing legal frameworks for the digital age and encourage voluntary industry initiatives around transparency, accountability and security.

AI is not a force to be feared, but a tool to be harnessed wisely in service of democratic values and aspirations. With the right governance frameworks and social norms in place, the United States can and must retain our global leadership in this critical technological domain. Ceding the AI race to less open societies would not only forfeit the profound

benefits for American consumers and businesses, but leave the future trajectory of this powerful technology in the hands of authoritarian regimes.

The choices we make today about how to approach AI governance will shape the fabric of American competitiveness, security and liberty for generations to come. We urge the NIST to reject reactive, heavy-handed proposals that would stymie our capacity to lead the AI revolution, and to instead work with Congress and industry stakeholders to advance pragmatic, forward-thinking solutions to maximize this technology's positive impact while mitigating its avoidable harms.

Thank you again for the opportunity to share our perspective on these critical issues. We look forward to continuing this important dialogue.

Sincerely,

Steve DelBianco
President and CEO
NetChoice