Before the
Office of Science and Technology Policy
Washington, D.C.

In re	Document Number: 2023-11346
Request for Information; National Priorities for Artificial Intelligence	

COMMENTS OF
THE COMPUTER & COMMUNICATIONS INDUSTRY ASSOCIATION (CCIA)

In response to the Request for Information ("RFI") published in the Federal Register at 88 Fed. Reg. 34194 (May 26, 2023), the Computer & Communications Industry Association ("CCIA")[1] submits the following comments to the Office of Science and Technology Policy ("OSTP") responding to selected questions.

Protecting rights, safety, and national security:

1. What specific measures – such as standards, regulations, investments, and improved trust and safety practices – are needed to ensure that AI systems are designed, developed, and deployed in a manner that protects people's rights and safety? Which specific entities should develop and implement these measures?

To ensure that AI systems are designed, developed, and deployed in a manner that protects people's rights and safety, targeted regulations must address privacy, bias, and accountability through a risk-based framework. High impact systems should be subject to more stringent regulation and scrutiny, while low impact systems may require limited guardrails or even no such measures.

Aspects of AI that are not unique to the technology should be governed by existing law with discrete additions in the limited instances where AI introduces unique challenges. This will result in a predictable and stable environment for AI investment, while limiting duplicative regulation and regulatory arbitrage.

While coordination is necessary, we do not recommend the creation of a new agency or department to develop or implement these measures. Particularly given how widely AI will affect our society, impacting many different sectors, a new agency will lack subject matter

[1] CCIA is an international, not-for-profit trade association representing a broad cross-section of communications and technology firms. For more than 50 years, CCIA has promoted open markets, open systems, and open networks. CCIA members employ more than 1.6 million workers, invest more than $100 billion in research and development, and contribute trillions of dollars in productivity to the global economy. A list of CCIA members is available at https://www.ccianet.org/members.

expertise and would likely lead to regulatory duplication and stifle investment in and development of AI systems. Instead, a coordinating role, housed within one of the following entities should coordinate the Federal Government's regulatory approach to artificial intelligence.

a. National Institute of Standards and Technology (NIST) or Similar Government Entity: NIST, with its expertise in coordinating technology policy across the government and its expertise in working with industry stakeholders to develop consensus approaches, can coordinate regulatory efforts and foster best practices across various domains. It can develop guidelines and standards for AI systems and facilitate collaboration among different agencies and industry stakeholders.

b. Other Existing Agencies: In most cases, existing agencies responsible for specific areas of law should oversee the regulation of AI falling within their jurisdiction. These agencies should collaborate with NIST and other relevant entities to ensure a coherent regulatory landscape. For example, the Federal Trade Commission (FTC) can enforce privacy and consumer protection regulations, the Food and Drug Administration (FDA) can address the usage of AI in medicine and health care, and the Equal Employment Opportunity Commission (EEOC) can address concerns related to AI bias in employment.

c. Industry Stakeholders: Industry stakeholders, including AI system developers, researchers, and organizations, should be actively consulted in the development of best practices and similar rules. Their expertise and insights are crucial for creating effective and practical regulations that align with industry realities.

2. How can the principles and practices for identifying and mitigating risks from AI, as outlined in the Blueprint for an AI Bill of Rights and the AI Risk Management Framework, be leveraged most effectively to tackle harms posed by the development and use of specific types of AI systems, such as large language models?

It's important to distinguish between different forms of artificial intelligence. Generative artificial intelligence and LLMs are only one form of AI, and regulating all forms of AI as if they were the same will result in negative impacts on American leadership in AI. However, the societal benefits ushered in by AI, such as improving weather forecasting and translating documents, hint at AI's tremendous potential if implemented responsibly with appropriate risk management. Because AI is so multifaceted and diverse, any proposed regulation should be flexible and employ a risk-based framework. The NIST AI Framework provides a good example of this approach. It empowers developers to find appropriate solutions while avoiding overly prescriptive principles that may stifle innovation.

Clear delineations should also be established to distinguish between AI as a contributing factor in decision-making and instances where AI makes decisions without human review. While AI can process and analyze data at an unprecedented scale, it may still lack contextual

understanding, ethical reasoning, and the ability to account for complex human values. Human review and judgment may be necessary to validate, interpret, and contextualize the outputs generated by AI systems. AI decisions made without human review may need to adhere to more detailed guidance and implement transparency measures to ensure fairness, non-discrimination, and compliance with legal requirements.

3. Are there forms of voluntary or mandatory oversight of AI systems that would help mitigate risk? Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?

Regulations should focus on identifying and addressing the concerns associated with AI development and deployment. The level of acceptable risk, required guardrails, and potential impacts should be evaluated based on the specific context. Lower impact applications may allow for higher tolerable risk levels and fewer guardrails. Higher impact applications permit less tolerated risk and require more robust guardrails.

Potential risks should be assessed on an ongoing basis throughout the AI system's lifetime. The Organization for Economic Cooperation and Development's ("OECD") AI Principles[2] illustrate a potential approach to creating these risk assessment models. We endorse the OECD's recommendation that AI systems should not pose unreasonable safety risks, as defined by existing laws and regulations in areas like consumer protection, in conditions of normal or foreseeable use or misuse throughout their lifecycle. Taking existing law and regulation as a starting point, and applying it to AI systems, will ensure that those systems reflect the legal guidelines companies must already abide by. At the same time, adaptation to reflect the unique aspects of AI may be helpful. This can include traceability to help to understand outcomes, to prevent future mistakes, and to improve the trustworthiness of the AI system, and a risk management approach to "identify, assess, prioritize and mitigate potential risks that can adversely affect a system's behavior and outcomes."

In some extremely high-risk arenas, a licensing regime might be appropriate, but such a regime would generally impede the development and deployment of AI systems and should be avoided in most applications.

For many applications, adherence to a set of industry best practices is likely to prove sufficient, and incentives or safe harbors could be created to promote such adherence. One potential model may be the Digital Trust & Safety Partnership[3] ("DTSP"), a new initiative of leading digital services dedicated to developing industry best practices to ensure consumer trust and safety that are verified through internal and independent third-party assessments. Critically,

---

[2] *OECD AI Principles Overview*, OECD.AI https://oecd.ai/en/ai-principles (last visited June 28, 2023).
[3] *Digital Trust & Safety Partnership Best Practices Framework*, https://dtspartnership.org/best-practices/ (last visited June 28, 2023).

these best practices were developed by trust and safety practitioners. Any best practices on AI should similarly be developed by practitioners and technical experts, not policymakers. Much like employing existing law and regulations, holding AI to the same industry standards as other technologies will help ensure that systems are designed, developed, and deployed inside the realm of acceptable and responsible use.

4. What are the national security benefits associated with AI? What can be done to maximize those benefits?

The United States has long been at the forefront of AI innovation, with a strong foundation driven by its vibrant tech industry, prestigious universities, and world-class research institutions. At a base level, retaining this leadership will allow the United States to maintain its competitive position globally.

Maintaining that leadership is critical to national security. There are legitimate concerns about the use of AI in government surveillance and policing; by ensuring AI development is led in the U.S., we can ensure that the development embeds democratic principles rather than being forced to rely on AI developed in nations that do not share those same principles. Domestic development also minimizes the risk of adversaries embedding backdoors into technology that the United States relies upon for its security. Furthermore, AI offers national security benefits such as improved situational awareness, enhanced intelligence and surveillance, cybersecurity, autonomous systems, and decision support. To maximize these benefits, governments should invest in research, develop AI talent, foster collaboration, ensure ethical use, and engage in cooperation with strategic allies.

5. How can AI, including large language models, be used to generate and maintain more secure software and hardware, including software code incorporating best practices in design, coding and post deployment vulnerabilities?

AI, including large language models, can enhance software and hardware security through various means. For instance, it can analyze code for vulnerabilities, automate security patching, detect anomalies and threats, provide guidance in secure design and coding, aid in vulnerability discovery and mitigation, and contribute to the development of adaptive and self-defending systems. Collaboration between AI researchers, security experts, and developers is vital to leverage AI effectively, ensuring alignment with best practices and continuous improvement through research and development. For example, in recent years, Meta began using AI-based code analysis tools to identify security issues[4] in its software and proactively address

[4] Jerome Pesenti, *Facebook's five pillars of Responsible AI*, Meta AI (June 22, 2021), https://ai.facebook.com/blog/facebooks-five-pillars-of-responsible-ai/.

them, while Google[5] employs AI algorithms to automate the identification and application of security patches across its software and systems.

6. How can AI rapidly identify cyber vulnerabilities in existing critical infrastructure systems and accelerate addressing them?

AI can rapidly identify cyber vulnerabilities in critical infrastructure systems by automating vulnerability scanning, real-time threat monitoring, predictive analytics, deep learning for anomaly detection, vulnerability prioritization, and automated patch management. By leveraging these AI capabilities, organizations can accelerate the identification and resolution of vulnerabilities, enabling proactive security measures and enhancing the resilience of critical infrastructure systems.

CCIA's members have engaged in responsible AI development, ranging from developing and applying their own responsible AI principles to conducting academic research that promotes privacy-by-design and the hardening of AI against motivated attackers seeking to extract training data, among other valuable contributions. In June 2023, Google[6] introduced its Secure AI Framework, which helps automate defenses to keep pace with existing and new threats and contextualize AI system risks among other measures. Google is working directly with organizations, including customers and governments, to help them understand how to assess AI security risks and mitigate them. Mitigation techniques include workshops with practitioners and publishing best practices for deploying AI systems securely. Similarly, Amazon's AWS[7] offers services to help AI systems developers better detect bias in datasets and models, provide insights into model predictions, and better monitor and review model predictions through automation and human oversight.

Advancing equity and strengthening civil rights:

9. What are the opportunities for AI to enhance equity and how can these be fostered? For example, what are the potential benefits for AI in enabling broadened prosperity, expanding economic and educational opportunity, increasing access to services, and advancing civil rights?

Transparency and disclosure are crucial aspects of designing thoughtful, adaptable regulation that can be applied in all contexts. Access to relevant information about how an AI system was designed, trained, and operates helps to ensure accountability and user trust. Individuals should be able to comprehend the foundations of decisions made by AI. At the same

---

[5] Royal Hansen & Phil Venables, *Responsible AI practices*, The Keyword (June 8, 2023), https://ai.google/responsibility/responsible-ai-practices/.
[6] *Id*.
[7] AWS, *Responsible Use of Artificial Intelligence and Machine Learning*, https://aws.amazon.com/machine-learning/responsible-machine-learning/ (last visited June 28, 2023).

time, the information provided must be relevant information and should take into account the specific context of the AI system. AI systems involved in significant decisions like those in medical or financial situations may require a higher level of disclosure, while systems operating in less significant areas—systems that automate product descriptions or flag posts for human review by a moderator, for example—may require little or even no disclosure.

That being said, broad agreement exists among leading AI developers and researchers, including CCIA's members, that responsible AI development requires the following:

- Design for social benefit.
- Design to avoid unfair outcomes.
- Analyze and minimize risks as you design.
- Consider the risks to third parties from AI systems during design, but also the benefits.
- Use up-to-date safety, security, and privacy best practices.
- Monitor and govern identified risks in deployed systems.
- Provide appropriate disclosures for deployed AI systems.

By engaging in this type of responsible AI development, AI systems can help to enhance equity.

10. What are the unique considerations for understanding the impacts of AI systems on underserved communities and particular groups, such as minors and people with disabilities? Are there additional considerations and safeguards that are important for preventing barriers to using these systems and protecting the rights and safety of these groups?

To avoid unfair outcomes, AI systems developers must take proactive measures such as collecting diverse and unbiased data, defining fairness metrics, and involving multidisciplinary teams. It is important to address biases in data, ensure transparency and explainability of AI systems, and implement regular auditing and monitoring for fairness. User feedback channels and redress mechanisms can provide avenues for addressing unfair outcomes, while regulatory compliance can enforce accountability and promote responsible AI practices. Existing laws concerning civil rights and discrimination must factor into the design and development of AI systems. AI systems should uphold existing rules of law while ensuring that new technologies do not create new obstacles.

AI systems can even eliminate old obstacles. As just one example of this, artificial intelligence tools that can create a description of an image can be used by visually impaired individuals in conjunction with a screen reader to better enable them to use existing networks. While approaches such as the inclusion of alt text can help the visually impaired, they are not universally used. AI image description systems can help to eliminate accessibility barriers of this kind, as well as many others.

11. How can the United States work with international partners, including low- and middle-income countries, to ensure that AI advances democratic values and to ensure that potential harms from AI do not disproportionately fall on global populations that have been historically underserved?

The United States can lead global efforts to ensure that AI development upholds democratic values and protects historically underserved populations by sharing expertise and providing technical assistance to low- and middle-income countries. For example, the United States can offer training programs and resources to help these countries develop AI capabilities so long as they adhere to democratic principles and responsible AI development practices.

Nations should collectively address biases and discriminatory practices in AI systems in a way that advances democratic values. By working together to develop tools and methodologies, countries can strive for fair and equitable AI outcomes. By aligning their standards, countries can protect the rights and safety of individuals globally through privacy protection, algorithmic transparency, and accountability. Working with low- and middle-income countries in the course of AI development will help ensure not just that harms do not fall on them, but that the AI that is produced is better by incorporating diverse perspectives and data. Furthermore, collaborating with these nations lessens the possibility that they will engage with geopolitical rivals who do not have the same commitment to democracy, harming the development of their democratic institutions.

Through these collective efforts, with the United States at the forefront, AI can be developed and deployed in a way that advances democratic values and avoids disproportionate harm to historically underserved populations. By promoting inclusivity, fairness, and respect for human rights on a global scale, the benefits of AI can be shared more equitably, ensuring a more inclusive and just AI ecosystem.

12. What additional considerations or measures are needed to assure that AI mitigates algorithmic discrimination, advances equal opportunity, and promotes positive outcomes for all, especially when developed and used in specific domains (e.g., in health and human services, in hiring and employment practices, in transportation)?

As a recent letter from the FTC, CFPB, DOJ Civil Rights Division, and EEOC[8] noted, these agencies already have the legal ability to enforce anti-discrimination law with regard to AI systems. AI systems must be designed to adhere to existing laws, including those that bar discrimination. Ethical design and development with careful consideration of existing discrimination law is key to creating technologies that maximize benefit while minimizing detriment. For example, AI algorithms used for job candidate screening and selection must align

[8] *FTC Chair Khan and Officials from DOJ, CFPB and EEOC Release Joint Statement on AI* (April 25, 2023), https://www.ftc.gov/news-events/news/press-releases/2023/04/ftc-chair-khan-officials-doj-cfpb-eeoc-release-joint-statement-ai.

with anti-discrimination laws, ensuring fair opportunities for all candidates. AI-driven systems for credit scoring and loan approvals should comply with laws that prohibit discrimination in lending, promoting transparency and fairness in decision-making.

Again, discrimination is already illegal and this applies to both existing and new AI technology. Any AI technology that does not follow existing law should not be deployed, and regulators can and should enforce existing law against such an AI technology if it is deployed.

CCIA would caution against AI-specific provisions that set a different standard for AI systems and humans, as the difference in standard between the two could and would be exploited to the detriment of underserved populations.

13. How might existing laws and policies be updated to account for inequitable impacts from AI systems? For example, how might existing laws and policies be updated to account for the use of generative AI to create and disseminate non-consensual, sexualized content? Bolstering democracy and civic participation?

To address the inequitable impacts of AI systems, including the use of generative AI for creating and disseminating non-consensual, sexualized content, and to bolster democracy and civic participation, existing laws can be applied with minimal additions to account for AI-specific situations if necessary.

As noted above in response to Question 12, there is no need for an update to say "racial discrimination by an AI is illegal." Discrimination is already illegal. Amendments may be required to address issues of intent, as AI systems do not form an intent, and liability, to determine what entity bears responsibility, but broadly speaking, oversight should target outcomes, rather than mechanisms. Any amendment should apply to non-AI and AI applications alike, ensuring that the law is technology-neutral and focuses on protecting individuals from harm. It should not matter whether discrimination is done by an AI system or a human, or whether a non-consensual image is produced by AI or by a human; in either case, the conduct is what is harmful and should be prevented, regardless of the actor.

Analogies from product liability law and from contributory infringement law may prove helpful in assessing the proper liable party. For example, where a system is suitable for many uses but could potentially be abused in an illegal way, liability should not lie on the system's creator but rather on the individual actually abusing it. At the same time, if the system was designed such that it is specially adapted to be used in violation of the law, liability on the system's designer might be appropriate.

14. How can AI be used to strengthen civic engagement and improve interactions between people and their government?

AI tools can provide numerous positive opportunities. For example, AI tools can be leveraged to enhance civic engagement and government interactions by providing citizens with initial information about government processes and enabling the government to allocate resources more effectively. Through AI-powered systems such as chatbots and virtual assistants, citizens can access relevant information and navigate government services more efficiently. These tools can also help improve transparency and empower citizens to stay informed and engaged. Additionally, AI's predictive analytics capabilities enable governments to identify areas that would most benefit citizens, optimizing resource allocation and enhancing the delivery of public services. By harnessing the power of AI, governments can better serve their constituents and foster a more engaged and participatory democracy.

15. What are the key challenges posed to democracy by AI systems? How should the United States address the challenges that AI-generated content poses to the information ecosystem, education, electoral process, participatory policymaking, and other key aspects of democracy?

Understanding the scope of these challenges through comprehensive research is crucial to addressing them effectively. All AI systems are already required to follow existing law. Where existing law has anticipated challenges to democracy, enforcement against AI systems can address these expected concerns. But to address challenges specific to AI, we must understand its nuances, and supplement existing law with equally specific amendments.

One key challenge is the spread of AI-generated content, including realistic deepfakes and disinformation, which can undermine the integrity of the information environment by making it difficult for users to discern between real and manipulated sources. For instance, AI-generated deepfake videos could be used to spread false narratives about political candidates, potentially influencing elections. While this possibility poses risks to public trust, informed decision-making, and the credibility of democratic processes, it is not a new phenomenon and should not be regarded as such. Doctored photos were employed long before generative AI became available.[9] AI disinformation should be viewed through the lens of other disinformation and specially regulated only to the extent that it diverges in ways that require unique regulation.

However, there are some key challenges we can anticipate and address preemptively, such as potential bias and discrimination that pose risks to equal opportunity and democratic principles. If AI design and development adheres to existing anti-discrimination laws, it would be more difficult for the resulting AI system to perpetuate and amplify existing societal biases. It is better to regard AI as one part of a cohesive information ecosystem that abides by democratic

[9] Drew Harwell, *Doctored images have become a fact of life for political campaigns. When they're disproved, believers 'just don't care.'*, Washington Post (Jan. 14, 2020); https://www.washingtonpost.com/technology/2020/01/14/doctored-political-images/.

principles, than to allow it to develop independently outside of those guardrails while it awaits specialized instruction.

Innovating in public services:

24. How can the Federal Government effectively and responsibly leverage AI to improve Federal services and missions? What are the highest priority and most cost-effective ways to do so?

AI will have widespread impact, touching many missions of the Federal Government, and its application should be explored where it is feasible and appropriate. The Federal Government can effectively and responsibly leverage AI to enhance Federal services and missions by prioritizing its adoption in areas where it can deliver better outcomes using existing resources.

Like the Federal Government's response to the growth of computing and resulting productivity gains, it can harness AI to achieve similar advancements in efficiency, effectiveness, and innovation. By strategically integrating AI into various processes and systems, existing agencies can streamline operations, automate repetitive tasks, improve decision-making, and deliver services more efficiently to the public. As one example, AI fire and flood forecasting tools could be used by the Federal Government as part of its disaster response capabilities.

To maximize the benefits of AI adoption, the Federal Government should prioritize areas with the highest potential for impact and cost-effectiveness. This includes leveraging AI in data analysis and insights generation, predictive modeling, resource allocation, fraud detection, cybersecurity, and personalized services. By focusing on these priority areas, the Government can optimize resource utilization, enhance service delivery, and achieve improved outcomes across a range of Federal missions.

Furthermore, a responsible approach to leveraging AI involves ensuring transparency, accountability, and fairness in AI systems. Clear guidelines, standards, and ethical frameworks should be established to guide the development, deployment, and oversight of AI applications within the Federal Government. Additionally, fostering partnerships with industry, academia, and other stakeholders can facilitate knowledge sharing, collaboration, and the adoption of best practices in AI implementation.

By recognizing the broad impact of AI and strategically deploying it in priority areas, the Federal Government can unlock significant improvements in Federal services and missions, ultimately enhancing efficiency, productivity, and public value.

25. How can Federal agencies use shared pools of resources, expertise, and lessons learned to better leverage AI in government?

The Federal Government should not establish an agency specifically equipped to deal with artificial intelligence, but integrate AI tools and resources throughout existing agencies that wield the expertise to optimize it. By tapping into the expertise and experiences of different agencies, the Federal Government can maximize the benefits of AI adoption and drive positive outcomes across various mission areas.

By adopting this approach, the Federal Government can avoid duplication of efforts, promote collaboration, and effectively harness the diverse knowledge and experience within different agencies. Shared resources can include AI infrastructure, data repositories, and computational resources, which can be made accessible to multiple agencies for AI initiatives. This approach encourages efficiency and cost-effectiveness by maximizing the utilization of existing resources.

Furthermore, expertise and lessons learned from one agency can be shared with others, facilitating knowledge transfer and accelerating AI adoption across the government. Collaboration can occur through interagency working groups, knowledge-sharing platforms, and partnerships with external stakeholders, such as academic institutions and industry experts. This collaborative ecosystem enables agencies to leverage each other's experiences, successes, and challenges, leading to more informed decision-making and avoiding redundant efforts.

Integrating AI governance throughout existing agencies ensures that expertise is embedded in relevant mission areas. Such integration also allows agencies to tailor AI implementation to their specific needs and goals while ensuring compliance with legal, ethical, and regulatory requirements. This distributed approach enables agencies to leverage their domain knowledge and apply AI in ways that optimize outcomes and address the unique challenges of their respective missions.

28. What can state, Tribal, local, and territorial governments do to effectively and responsibly leverage AI to improve their public services, and what can the Federal Government do to support this work?

Within the past year several states have begun to more closely analyze the ways in which artificial intelligence is currently being used, as well as opportunities to further leverage AI technologies. These efforts have primarily been channeled through the creation of State Task Forces whose goals are to examine the various ways in which state agencies are utilizing AI tools currently and their impact.[10]

[10] *See, e.g.,* Connecticut SB01103; Illinois HB148734; Texas HB2060; Louisiana SCR49; New York S. 6402.

We applaud these efforts by states to further study AI, particularly as these task forces aim to ensure that AI tools are being used responsibly when involved with sensitive decisions, such as decisions impacting benefits or housing. We also encourage states to establish a designated coordinating role for AI, a model Vermont took when it appointed an "AI Director" serving under the CIO. Establishing such a role can help streamline efficiency across state agencies, as the coordinator can establish templates and guidelines for the use of AI tools across various agency activities, while also ensuring that innovation is not being inhibited and can be responsible for ensuring that the state is staying up-to-date with the latest AI technologies.

As states wait for potential federal-level regulation of AI tools, it is important to ensure that interoperability and consistency are central to the development of any policy in order to avoid creating 50 different state frameworks, which would confuse both consumers and businesses. To that end, any approach states take should be principled and risk-based to avoid potentially deterring any efforts aimed at beneficial AI use development. Further, states should consider existing laws and frameworks at their disposal that can mitigate concerns surrounding AI and its impacts, such as state unfair and deceptive practices laws in addition to existing laws regarding discrimination.

* * *

CCIA appreciates the opportunity to provide input on this important issue. We look forward to working with OSTP as the Administration's approach to AI governance continues to develop and would be happy to provide additional information and resources.

Respectfully submitted,

Joshua Landau
Senior Counsel, Innovation Policy
Erin Sakalis
Law Clerk
Computer & Communications Industry Association
25 Massachusetts Ave NW
Suite 300C
Washington, DC 20001
jlandau@ccianet.org