Policy Recommendations for the Responsible Use of Artificial Intelligence

June 2024

Data & Trust Alliance

As regulators and policymakers around the world try to understand, define and control Artificial Intelligence, Machine Learning and Generative Artificial Intelligence, the number of approaches, methods and requirements are varied and challenging for business enterprises to navigate as they continue to innovate and develop new products and services. The simultaneous evolution of regulatory requirements and innovations creates an environment that can be difficult to navigate even when most organizations and governments have the same goals in mind: to build a future that creates net benefits for individuals, communities and society as a whole. To navigate this challenging moment in time takes dedication, focus on governance and a commitment to trusted data and analytic practices.

The Data & Trust Alliance, an organization of world-class business enterprises, created this policy recommendation document as a first step in providing guidance to all stakeholders as we attempt to create a positive future that effectively accesses these new technologies. We hope these recommendations are constructive as society and innovation progress and we look forward to examining these and other topics in more detail.

Published 06.13.24

Authored by the AI Policy Working Group of the Data & Trust Alliance

❶ Regulate AI risk, not AI algorithms.

We believe that AI should be explainable, fair, robust, transparent, and privacy protective in a manner commensurate with the risk of the intended uses and purposes of the AI system. This assertion means that we:

1A

Support an approach that regulates use of AI in high-risk applications.

Regulation should encourage innovation while mitigating risk, based on the intended use-cases and purposes, taking into account factors such as the application(s), end-user(s), how reliant the end-user would be on the technology, and the level of human oversight. Accordingly, we support policy that would:

- Provide explicit categories with sets of high-risk AI use cases for which regulations would apply in order to provide clarity and predictability to AI developers, deployers and the public;
- Mandate impact assessments and bias testing for high-risk AI use cases;
- Require transparency, defined by users knowing when they are interacting with an AI system within a high-risk use case and whether they have recourse to engage with a real person, if desired. Also, encourage explainability, where appropriate and commensurate

 © 2024 The Data & Trust Alliance, all rights reserved

with the risk, to give society better visibility and greater assurances into how these models operate.

- Where appropriate and commensurate to the risk, AI developers should be required to disclose technical information about the development and performance of an AI model related to the use case(s) in order for deployers to conduct the requisite impact assessments and bias testing, as well as the empirical data sources used to train it, to give society better visibility and greater assurances into these models;

- Prevent and stop harm. Take measures so that AI systems are not leveraged for specific prohibited uses that present a significant risk of harm. These prohibitions include the use of AI for mass surveillance, racial profiling, and violations of basic human rights and freedoms;

- Ensure AI developers and deployers employ "trustworthy" AI governance including safety, privacy, disclosure, data quality, etc., which may include, where appropriate, leveraging the NIST AI Risk Management Framework and its best practices around these trustworthy AI qualities and/or the White House Voluntary AI Commitments for Generative AI. Good AI governance should be calibrated for the specific risks and use cases, and may include:

  → Built and tested for safety; continue to develop and apply strong security practices to avoid unintended results that create harm.

  → Incorporate privacy design principles, give opportunity for notice and consent for consumers.

  → Enable appropriate disclosure when personal information is processed for automated decision-making or profiling impacting matters of significance, including options to elect human review or alternative processing.

  → Development of cross-sector data provenance standards to ensure trustworthy data and to reject sources of data, deemed to be untrustworthy (D&TA effort underway).

Please see included addendum, "Notable AI, Cybersecurity, and Privacy Commitments by D&TA Members" (individually and via association groups).

Learn more about the Data & Trust Alliance

1B

Avoid government "AI license to operate" obligations

Mandatory government licensing and pre-market deployment certification and requirements—particularly when applied to non-high risk AI systems—have significant economic costs, stifle competition, and reduce the availability of opensource AI systems. Policymakers should recognize:

- AI value chains are complex and constantly evolving. Requiring a license from the government at any point in this value chain would create an enormous obstacle to its efficient operation.

© 2024 The Data & Trust Alliance, all rights reserved

- The White House Voluntary AI Commitments for frontier models should be promoted and any new policies should align to and reinforce them [rather than create government licensing obligations].

- Government licensing requirements for AI would pose a significant obstacle to the development and adoption of AI systems, including open source. With thousands of start-ups and open source contributors working to make AI systems better, the open source community can more quickly find potential risks in systems and work to mitigate them. Open source AI also enhances innovation and competition by enabling small businesses and startups to focus on building new and innovative products, rather than paying to access proprietary models.

1C

Ensure any efforts to regulate general purpose AI are proportionate and technically feasible.

Recent advancements in the use of foundation models—AI models trained to perform general functions, such as text or image generation or recognition, rather than for a specific purpose—are important building blocks in a flexible, robust, and innovative AI ecosystem. The adaptability of these models enables them to be used across a range of applications; the flexibility also means that bad actors can misuse them. Their flexibility, however, does not warrant a fundamental departure from a risk-based approach. We should formalize best practices around the evaluation and disclosure of foundation models/general purpose AI to help in the responsible deployment of AI. Regulation should strive to be commensurate with the risks associated with the use of technology and the foreseeable uses of technology.

❷ Use existing sector-specific regulatory authorities which are best able to regulate AI use, and use supported and effective existing regulations.

Policymakers should recognize that:

- AI is being deployed in numerous fields where there are already regulations and regulators that can ensure the safe deployment of these systems (e.g., health care, education, financial services, etc.). We already have regulations and regulators with oversight of decision-making in specific domains (e.g., DOT for vehicles, FDA for medical devices, the PTO and Copyright Office for intellectual property), who are regulating AI use-cases already and/or are well-positioned to address domain-specific concerns of AI going forward.

- We should leverage these existing structures, where possible, rather than creating a new regulator with an overlapping zone of responsibility.

- We recognize that there may be potential gaps in the current legal structure that should be considered. Governments should conduct a comprehensive assessment of where gaps, such as the lack of a comprehensive federal privacy legislation, might occur in existing regulatory structures before creating new agencies or broadening the powers of existing agencies beyond their subject matter expertise.

© 2024 The Data & Trust Alliance, all rights reserved

→ Where possible, policymakers should strive for consistency and harmonization in definitions and frameworks in order to reduce frictions and enhance the innovative potential of AI systems, particularly for low-risk use cases.

Instead, ensure each agency can appropriately address AI risks within their areas of expertise and statutory authority allowing an agile, collaborative and consistent approach to AI governance. Additionally, ensure that each agency has a clearly defined scope of authority to avoid regulatory overreach. Ensure existing regulators have the knowledge and resources necessary to mitigate the risks of AI in their domains. Existing regulators are best able to manage risk. Embed AI rules and guidelines into existing frameworks, as appropriate.

History has shown that agencies evolve to meet the challenges of modern technology. NIST's AI Risk Management Framework is a great example. NIST was founded over 100 years ago, yet it was capable of producing its RMF, which will serve as the foundation for AI risk management in industry, as well as inform regulation.

❸ Differentiate the compliance responsibilities of developers and deployers, including data collection and usage practices.

We support trustworthy AI and policies that promote it by establishing clear, risk-based guardrails and liabilities, which are tailored to the roles and capabilities individual organizations play in the broader AI developmental lifecycle. When crafting more specific rules and regulations, policymakers should be aware of the many differences—in capabilities and responsibilities—that exist within these broad and diverse ecosystems and in supply chains that require complex coordination and compliance obligations. These responsibilities need to be further explored as practices and innovation continue to develop, recognizing intellectual property rights and principles such as fair use.

❹ Mitigate risks by investing in R&D, education, and workforce development.

Addressing the risks AI poses will require a combination of smart policy, education, and good science. A major challenge limiting progress on how effectively we can study AI risks is limited access to computing power. We need to have an open, inclusive ecosystem of researchers and diverse stakeholders developing and evaluating AI models.

We support the recommendations of the National AI Research Resource Task Force to share computing resources to help create the infrastructure necessary to foster a diverse, open, inclusive R&D ecosystem. There are a wide range of R&D targets that policymakers and research agencies should be pursuing, including watermarking and synthetic media detection techniques and better bias testing methodologies for multi-modal models. The recently released National AI R&D Strategic Plan provides a good outline of the areas on which we should be focusing our efforts.

We also support increasing investment in AI education at all levels to better prepare students for future AI-related job opportunities and to encourage consumer fluency with AI systems to enhance trust in those systems and to encourage critical evaluation of AI outputs.

© 2024 The Data & Trust Alliance, all rights reserved

This education would involve building AI programming into educational curricula and directing greater research funding for AI testbeds towards minority-serving institutions, ensuring a more diverse range of stakeholders guide the design, development, and application of AI systems. While the AI field today may not reflect the demographics of our society, moving towards a more diverse AI ecosystem—from developers to users—could enable organizations to avoid and mitigate unwanted AI bias by including and reflecting the interests and values of communities likely to be impacted.

Next Steps

Given the significant policy and regulatory activity on artificial intelligence, both domestically and internationally, in the last month, the D&TA Working Group hopes this document is a helpful departure point for critical discussion on the key issues impacting this cross-sector group of industry leaders.

Authored and endorsed by a subset of Data & Trust Alliance member organizations

- AARP
- GM
- Howso
- Humana
- IBM
- Mastercard
- Meta
- NFL
- Nielsen
- Nike
- Transcarent
- Walmart
- Warby Parker

Notable AI, Cybersecurity and Privacy Commitments by Data & Trust Alliance Members (individually and via association groups)

- **Automakers' Commitment to Privacy: Consumer Privacy Protection Principles for Vehicle Technologies and Services**
- **Automotive Cybersecurity Best Practices**
- **Consumer Technology Association Principles for the Privacy of Personal Health Data**
- **US Chamber of Commerce Privacy Principles**
- **IBM's Approach to AI Ethics**
- **Business Roundtable Policy Perspectives on AI**
- **Partnership on Responsible AI**
- **NIST AI Risk Management Framework**
- **Deloitte AI Risk Management Framework**
- **Meta Responsible AI Pillars**
- **CVS Health Privacy Center**
- **Nielsen Privacy Principles**
- **AutoISAC Best Practices for Function-Based Approaches to Managing Cyber Risk**
- **White House Voluntary Commitments for Generative AI**
- **Mastercard Data Responsibility Principles**
- **Walmart Responsible AI Pledge**

About D&TA

The Data & Trust Alliance is a group of industry-leading enterprises committed to a future powered by the responsible use of data and AI. We leverage the collective expertise and influence of our members—among them the leading deployers of data and AI in business—to create and adopt practices that enhance trust in data, in AI models, in the people and process through which they are deployed. Our only KPI is adoption by practitioners. Learn more at dataandtrustalliance.org.

Please note that the materials do not necessarily reflect the views of any D&TA member organization. © 2024 The Data & Trust Alliance, all rights reserved