

March 25, 2024

The Honorable Gina Raimondo
Secretary
Department of Commerce
1401 Constitution Ave. NW
Washington, DC 20230

RE: Openness and Transparency in AI Provide Significant Benefits for Society

Dear Secretary Raimondo,

We, the undersigned civil society organizations and academic researchers, write to underscore key points of consensus about the importance of openness and transparency in AI models.

We applaud the Administration for its significant actions aimed at harnessing the benefits and mitigating the risks from AI across all sectors and domains.¹ We appreciate the opportunity to provide comments to the National Telecommunications and Information Administration's request for comment on openness in AI models, and many of our organizations are submitting more detailed comments in response.² An opportunity for public comment will be similarly vital if the Bureau of Industry and Security proposes export controls on AI models,³ which could have significant drawbacks for economic growth, democratic values, and people's safety.

Although we approach openness and transparency in AI from a wide range of perspectives, we all agree that it has a vital role to play in making AI worthy of our trust. We send this letter in order to underscore three broad points of consensus about openness and transparency in AI:

- 1. Open models can provide significant benefits to society, and policy should sustain and expand these benefits.** For decades, open source software has provided building blocks for everything from creating art to designing vaccines. According to recent estimates, open source software is worth more than \$8 trillion in value⁴ and is a part of 96% of commercial software.⁵ The U.S. government is one of the biggest users of

¹ See, e.g., President Biden, "[Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#)," The White House, October 2023.

² National Telecommunications and Information Administration, "[Dual Use Foundation Artificial Intelligence Models with Widely Available Model Weights](#)," Federal Register, February 26, 2024.

³ Alan Estevez, "[Fireside Chat with Under Secretary Alan Estevez](#)," Center for Security and Emerging Technology (CSET), Georgetown University, December 2023. ("We're talking about ... large language models, we're having those discussions ... I have a team ... working on what's the answer.") See also Karen Hao, "[The New AI Panic](#)," The Atlantic, October 2023. ("Commerce is considering a new blockade on a broad category of general-purpose AI programs, not just physical parts, according to people familiar with the matter.")

⁴ Manuel Hoffman et al., "[The Value of Open Source Software](#)," Harvard Business School, January 2024.

⁵ Synopsys, "[2024 Open Source Security and Risk Analysis Report](#)," February 2024. (Analyzed 1,067 commercial codebases across 17 industries in 2023, and found that 96% of those codebases contained open source.) See also, Chinmayi Sharma, "[Tragedy of the Digital Commons](#)," North Carolina Law Review, October 2022. ("Google, iPhones, the national power grid, surgical operating rooms, baby monitors, surveillance technology, and wastewater management systems all run on open-source software... Without it, our critical infrastructure would crumble.")

open source software in the world,⁶ and funds open source approaches ranging from boosting cybersecurity to protecting human rights and fighting cancer.⁷

Openness in AI can provide similar benefits. Indeed, many of AI's most promising applications have already been fueled by open source and open science,⁸ and openness can support key societal goals, such as:

- **Advancing innovation, competition, and research:** Open models promote economic growth by lowering the barrier for innovators, startups, and small businesses from more diverse communities to build and use AI. Open models also help accelerate scientific research because they can be less expensive, easier to fine-tune, and supportive of reproducible research.
- **Protecting civil rights and human rights:** Open models make it easier for regulators and civil society to assess AI systems for compliance with laws protecting civil rights, privacy, consumers, and workers. They increase transparency, education, testing, and trust around the use of AI, enabling researchers and journalists to audit and write about AI systems' impact on different demographic groups.⁹ And, they also lower the barrier for stakeholders outside of large tech companies to shape the future of AI, enabling more AI services to be built by and for diverse communities with different needs that big companies may not always address.
- **Ensuring safety and security:** Open models advance safety and security by accelerating our understanding of AI capabilities, risks, and harms through independent research, collaboration, and knowledge sharing. In turn, this supports regulators and researchers who need the latest methods, tools, and understanding to effectively monitor and test large scale AI systems.

2. Policy should be based on clear evidence of marginal risks that open models pose compared to closed models. Recent research outlines the importance of evaluating the risks of open models not in a vacuum, but in comparison to the risks and benefits from closed models and pre-existing technologies like the internet.¹⁰ Put another way, what is the marginal risk of an open model? For example, the claim that open models make it easier to operate disinformation campaigns needs to be compared against the ease of conducting disinformation campaigns using closed models like DALL-E 3 and

⁶ Eric Goldstein and Camille Stewart Gloster, "[We Want Your Input to Help Secure Open Source Software](#)," Cybersecurity and Infrastructure Security Agency, August 2023. See also, federal policy supporting open source and open innovation, e.g., Tony Scott and Anne Rung, "[M-16-21 Federal Source Code Policy: Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software](#)," August 2016.

⁷ See, e.g., Rachel Berkowitz, "[How Berkeley Lab Helped Develop One of the World's Most Popular Open-Source Security Monitoring Platforms](#)," Lawrence Berkeley National Laboratory, February 2023; "[Supporting Critical Open-Source Technologies That Enable a Free and Open Internet](#)," State Department, November 2023; and "[CANcer Distributed Learning Environment](#)," National Cancer Institute, February 2023.

⁸ E.g., Key model architectures like AlexNet, frameworks like PyTorch and TensorFlow, and research on topics like attention mechanisms were all made widely available, fueling significant advances in AI R&D.

⁹ See, e.g., Stephen Casper et al., "[Black-Box Access is Insufficient for Rigorous AI Audits](#)," arXiv, January 2024. ("[W]hite-box access to the system's inner workings (e.g., weights, activations, gradients) allows an auditor to perform stronger attacks, more thoroughly interpret models, and conduct fine-tuning.")

¹⁰ Sayash Kapoor et al., "[On the Societal Impact of Open Foundation Models](#)," Center for Research on Foundation Models (CRFM), Stanford University, February 2024.

existing tools like Photoshop.¹¹ Meanwhile, open models can often provide significant marginal benefits compared to closed models, as outlined above. We urge you to be rigorous in evaluating and targeting the specific risks from openness in AI, including developing better proxies for risk that are not solely based on the amount of computing power used to train a model.¹²

- 3. Policy should consider a wide range of solutions to address well-defined marginal risks in a tailored fashion.** We do not claim that openness is always beneficial, and there are some situations where openness may exacerbate risks from AI. However, heavy-handed approaches to restrict the availability of model weights, such as broad export controls, could come with significant negative consequences,¹³ may be impractical¹⁴ and may unconstitutionally hinder scientific dialogue.¹⁵ Enforcing specific areas of law to address particular harms, for example in the realms of civil rights and unfair trade, is poised to be more effective and less damaging than broad restrictions on general purpose software.¹⁶

We encourage you to coordinate closely with other agencies and White House components that have equities on this topic. We urge you and the rest of the Administration to support more R&D into open approaches for AI, and to work with the open source community to advance better standards for testing and releasing open models. We also urge you to ensure that NTIA's forthcoming report, as well as any decision to use export controls for AI models, goes through a robust interagency process that includes the agencies with responsibility for competition policy, civil rights, and scientific research — not just the agencies that oversee national security.

¹¹ See, e.g., Sayash Kapoor and Arvind Narayanan, "[How to Prepare for the Deluge of Generative AI on Social Media](#)," Knight First Amendment Institute at Columbia University, June 2023. ("[T]he bottleneck for successful disinformation operations is not the cost of creating it.")

¹² Rishi Bommasani, "[Drawing Lines: Tiers for Foundation Models](#)," CRFM, November 2023. ("the relationship between compute and impact is quite tenuous and not evidenced... there is no demonstration that compute robustly predicts results on risk evaluations, let alone demonstrations that compute predicts the impact foundation models have in society... compute is a measure of upstream resource expenditure, naturally divorced from downstream societal impact.")

¹³ For example, they could significantly restrict American innovation and economic growth in AI, much like broad export controls on encryption in early web browsers were a key inhibitor of international e-commerce. They could also restrict testing for safety, which often relies on access to open models. And, they could reduce competition, as the associated licensing regime could disproportionately harm small firms.

¹⁴ See, e.g., Carrick Flynn, "[Recommendations on Export Controls for AI](#)," CSET, February 2020. ("New export control regulations on general purpose AI software ... are unlikely to succeed and should not be implemented.")

¹⁵ E.g., courts may find that export controls on the publication of model weights implicate the First Amendment and that the government has not met its burden to justify the restriction on scientific speech, much as courts have previously held in regard to the publication of encryption software source code. See *Junger v. Daley*, 209 F.3d 481, 485 (6th Cir. 2000) ("Because computer source code is an expressive means for the exchange of information and ideas about computer programming, we hold that it is protected by the First Amendment."); *Bernstein v. U.S. Dep't of Just.*, 176 F.3d 1132, 1141 (9th Cir. 1999), reh'g granted, opinion withdrawn, 192 F.3d 1308 (9th Cir. 1999) ("[W]e conclude that encryption software... must be viewed as expressive for First Amendment purposes, and thus is entitled to the protections of the prior restraint doctrine.")

¹⁶ See, e.g., Rishi Bommasani et al., "[Considerations for Governing Open Foundation Models](#)," Stanford Institute for Human-Centered AI, December 2023. ("As with many other threat vectors, the best policy choke points may hence lie downstream. For example, the U.S. AI Executive Order aims to strengthen customer screening for purchasers of biological sequences.")

Thank you for your attention to these matters. For any questions or further discussion, please contact Kevin Bankston, Senior Advisor on AI Governance, Center for Democracy & Technology (kbankston@cdt.org) and Jennifer Hodges, Head of US Public Policy & Government Relations, Mozilla (jhodges@mozilla.com).

Respectfully,

Organizations

Accountable Tech
Allen Institute for Artificial Intelligence
Center for Democracy & Technology
Chamber of Progress
Computing Research Association
Creative Commons
Data & Society
Electronic Frontier Foundation
EleutherAI
Engine
Federation of American Scientists
Fight for the Future
Government Information Watch
Information Technology and Innovation Foundation
Kapor Center
Library Futures
Mozilla
National Fair Housing Alliance
New America's Open Technology Institute
Open Source Initiative
Partnership on AI
Public Knowledge
R Street Institute

Individual Academic Signers

Sayed Choudhury, Carnegie Mellon University
Michelle De Mooy, Georgetown University
Oren Etzioni, University of Washington
Ali Farhadi, University of Washington
Camille François, Columbia University
Shubha Ghosh, Syracuse University
Peter Henderson, Princeton University
Daniel E. Ho, Stanford University
Sayash Kapoor, Princeton University
Kevin Klyman, Stanford University
Anne Lambright, Carnegie Mellon University

Mark A. Lemley, Stanford University
David S. Levine, Elon University
Percy Liang, Stanford University
Daniel W. Linna Jr., Northwestern University
Meredith Martin, Princeton University
Arvind Narayanan, Princeton University
Joelle Pineau, McGill University
Nathan Reitinger, University of Maryland
Bruce Schneier, Harvard University
Dawn Song, UC Berkeley
Suresh Venkatasubramanian, Brown University
Keith Webster, Carnegie Mellon University
Kevin Werbach, University of Pennsylvania

CC:

Alan Estevez, Under Secretary of Commerce for Industry and Security
Laurie Locascio, Under Secretary of Commerce for Standards and Technology
Alan Davidson, Assistant Secretary of Commerce for Communications and Information
Elizabeth Kelly, Director of the U.S. AI Safety Institute
Saif Khan, Senior Advisor to the Secretary of Commerce

Jeffrey Zients, Assistant to the President and White House Chief of Staff
Arati Prabhakar, Assistant to the President for Science and Technology Policy
Jake Sullivan, Assistant to the President for National Security Affairs
Bruce Reed, Assistant to the President and White House Deputy Chief of Staff
Lorraine Voles, Assistant to the President and Chief of Staff to Vice President Harris
Lael Brainard, Assistant to the President for Economic Policy
Neera Tanden, Assistant to the President for Domestic Policy
Deirdre K. Mulligan, Principal Deputy U.S. Chief Technology Officer
Sethuraman Panchanathan, Director of the National Science Foundation
Rohit Chopra, Director of the Consumer Financial Protection Bureau
Lina Khan, Chair of the Federal Trade Commission
Charlotte Burrows, Chair of the Equal Employment Opportunity Commission
Nathaniel C. Fick, Ambassador at Large for Cyberspace and Digital Policy
Ben Buchanan, White House Special Advisor on AI
Helena Fu, Director, Office of Critical and Emerging Technologies, Department of Energy
Craig Martell, Chief Digital and AI Officer, Department of Defense
Jonathan Mayer, Chief S&T Advisor and Chief AI Officer, Department of Justice
Eric Hysen, Chief Information Officer and Chief AI Officer, Department of Homeland Security