March 27, 2024

Mr. Travis Hall
National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Ave NW
Washington, DC 20230

RE: Comments of ACT | The App Association to the National Telecommunications and Information Administration on Dual Use Foundation Artificial Intelligence Models With Widely Available Model Weights (Docket No. 240216-0052)

I. Introduction & Statement of Interest

ACT | The App Association (App Association) appreciates the opportunity to provide input to the National Telecommunications and Information Administration (NTIA) on on the potential risks, benefits, other implications, and appropriate policy and regulatory approaches to dual-use foundation artificial intelligence (AI) models for which the model weights are widely available.[1]

The App Association is a global trade association for small and medium-sized technology companies. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem that engage with verticals across every industry. We work with and for our members to promote a policy environment that rewards and inspires innovation while providing resources that help them raise capital, create jobs, and continue to build incredible technology. Today, the value of the ecosystem the App Association represents—which we call the app economy—is approximately $1.8 trillion and is responsible for 6.1 million American jobs, while serving as a key driver of the $8 trillion internet of things (IoT) revolution.[2] Alongside the world's rapid embrace of mobile technology, our members create the innovative solutions that utilize AI to power IoT across various modalities and segments of the economy.

From the App Association's perspective, AI is an evolving constellation of technologies that enable computers to simulate elements of human thinking, such as learning and reasoning. An encompassing term, AI entails a range of approaches and technologies, such as machine learning (ML), where algorithms use data, learn from it, and apply their

[1] https://www.federalregister.gov/documents/2024/02/26/2024-03763/dual-use-foundation-artificial-intelligence-models-with-widely-available-model-weights.

[2] ACT | The App Association, State of the U.S. App Economy: 2020 (7th Edition) (Apr. 2020), available at https://actonline.org/wp-content/uploads/2020-App-economy-Report.pdf

newly-learned lessons to make informed decisions, and deep learning, where an algorithm based on the way neurons and synapses in the brain change as they are exposed to new inputs allows for independent or assisted decision-making. AI-driven tools are having, and will continue to have, substantial direct and indirect effects on Americans. Some forms of AI are already being used to improve American consumers' lives today – for example, AI is used to detect financial and identity theft and to protect the communications networks upon which Americans rely against cybersecurity threats. Moving across use cases and sectors, AI has incredible potential to enable faster and better-informed decision making through cutting-edge distributed cloud computing. For example, healthcare treatments and patient outcomes stand poised to improve disease prevention and conditions, as well as efficiently and effectively treat diseases through automated analysis of x-rays and other medical imaging. From a governance perspective, AI solutions will derive greater insights from infrastructure and support efficient budgeting decisions. It is estimated that AI technological breakthroughs will represent a $126 billion market by 2025.[3]

As AI systems, powered by streams of data and advanced algorithms, continue to improve services and generate new business models, the fundamental transformation of economies across the globe will only accelerate. At the same time, AI's growing use raises a variety of challenges, and some new and unique considerations, for policymakers as well as those making AI operational today. The App Association appreciates the efforts of NTIA, and other federal agencies, to address AI safety, reliability, and innovation per the Executive Order Concerning Artificial Intelligence.

The App Association has worked proactively to develop consensus around AI governance and policy questions from across its diverse and innovative community of small businesses. As a result of these consensus-building efforts, the App Association has created comprehensive policy principles for AI governance,[4] which we append to this comment and urge NTIA (and other policymakers) to align with. Notably, the App Association's policy principles for AI governance and policy address quality assurance and oversight, recommending that any AI policy framework utilize risk-based approaches to ensure that the use of AI aligns with the recognized standards of safety, efficacy, and equity. Our AI policy principles also prioritize ensuring the appropriate distribution and mitigation of risk and liability by providing that those in the value chain with the ability to minimize risks based on their knowledge and ability should have appropriate incentives to do so.

The App Association appreciates NTIA's discussion of open model foundations in its request for information, and agrees that open model foundations can support competition and innovation, and further transparency. Models with widely available

[3] McKinsey Global Institute, *Artificial Intelligence: The Next Digital Frontier?* (June 2017), *available at* https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx.

[4] The App Association's Policy Principles for Artificial Intelligence are included in this comment as **Appendix A**.

weights can benefit from a feedback loop that includes users and developers, enabling eased feature improvements as well as identification and mitigation of risk, while spending less resources.

The App Association appreciates NTIA's requesting input on how today's open source software licensing approach could inform its approach to dual use foundation models. While not analogous (because open source software licenses do not encapsulate all components and capabilities of an AI model), open source licenses can be beneficial in many scenarios by harmonizing terminology, training, deployment, weights, and documentation/monitoring. However, such licenses cannot actively prevent malfeasance. Ultimately, the App Association believes that the market, not government, should organically develop open model licensing approaches.

Further, building on the above, we offer the following comments and recommendations to NTIA:

- ***Improve its categorization of foundation models:*** Categorizing foundation models as either "open" of "closed" will not reflect the important distinctions between key existing categories of foundation models. The degree of "openness" depends on a range of factors,[5] making the drawing of a hard line between "open" and "closed" arbitrary.

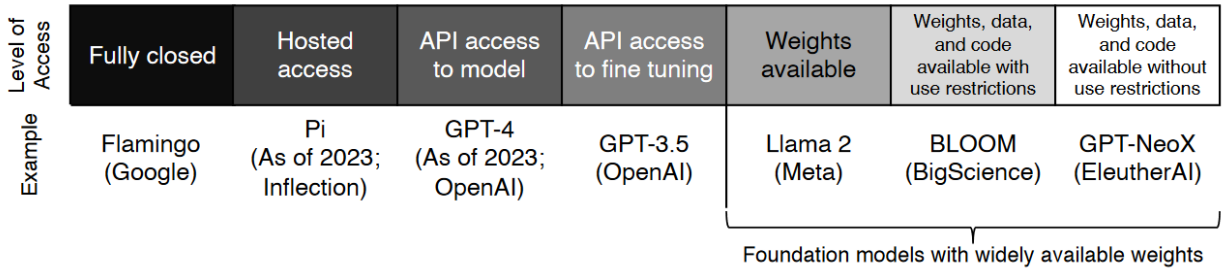

Bommasani, Rishi, et al. "Issue Brief Considerations for Governing Open Foundation Models | Stanford HAI." Stanford University Human-Centered Artificial Intelligence, 13 Dec. 2023, https://hai.stanford.edu/issue-brief-considerations-overning-open-foundation-models.

While each of the above categories of foundation model offers its own benefits and risks. While fully closed models may be preferrable to protect intellectual property, models that make weights available (or even source code) can provide access to a feedback loop with developers or the ability for users to make improvements.

For purposes of the Executive Order, we urge NTIA to ensure that a "dual-use foundation model" is not used synonymously with "open foundation model." The Executive Order provides the following:

[5] Bommasani, Rishi, et al. "Issue Brief Considerations for Governing Open Foundation Models | Stanford HAI." Stanford University Human-Centered Artificial Intelligence, 13 Dec. 2023, https://hai.stanford.edu/issue-brief-considerations-overning-open-foundation-models.

> "(k) The term "dual-use foundation model" means an AI model that is trained on broad data; generally uses self-supervision; contains at least tens of billions of parameters; is applicable across a wide range of contexts; and that exhibits, or could be easily modified to exhibit, high levels of performance at tasks that pose a serious risk to security, national economic security, national public health or safety, or any combination of those matters, such as by:
>
> > (i) substantially lowering the barrier of entry for non-experts to design, synthesize, acquire, or use chemical, biological, radiological, or nuclear (CBRN) weapons;
> >
> > (ii) enabling powerful offensive cyber operations through automated vulnerability discovery and exploitation against a wide range of potential targets of cyber attacks; or
> >
> > (iii) permitting the evasion of human control or oversight through means of deception or obfuscation.
>
> Models meet this definition even if they are provided to end users with technical safeguards that attempt to prevent users from taking advantage of the relevant unsafe capabilities."

The App Association urges NTIA to recognize that not all "open foundation models" reflect the characteristics described in the Executive Order, and to ensure that the scope of foundation models addressed in its report is confined to "dual-use foundation models" as defined in the Executive Order. If the scope of NTIA's report is not carefully restrained to this scope, it may cause definitional confusion in the short term, and later improperly expose foundation models that are not "dual-use foundation models" to future policy or regulatory requirements meant to be applied this category alone.

- ***Address harms that are demonstrable and systemic.*** For purposes of this exercise under the Executive Order, NTIA should focus on high-risk scenarios (e.g., health, safety) for which there is a clear evidence base to address (in other words, policy proposals should not be based on remote edge use cases or hypotheticals).
- ***Adhere to scalable risk-based harm mitigation principles.*** As NTIA explores policy and regulatory options for dual-use foundation models, we strongly urge NTIA to, consistent with the National Institute of Standards and Technology's AI Risk Management Framework, ensure that its proposals are grounded in utilizing risk-based approaches to ensure that levels of review, assurance, and oversight are proportionate to potential harms. Building on this foundation, NTIA should discourage blanket/one-size-fits-all approaches to risk mitigation for dual-use foundation models.

NTIA's definition of "widely available" should be similarly approached. The wide availability of a model is not necessarily an indicator of the risk(s) it may present. We urge NTIA's definition of "widely available" to reflect the harms presented by the relevant use case(s). Similarly, floating point operations do not necessarily indicate higher risks. Such definitional thresholds should primarily consider the capabilities of the model.

We urge NTIA to maintain a broad perspective in considering risk in this matter. Many other factors than weights can alter the risks and benefits for a foundation model, such as training data, evaluation metrics, and deployment guidelines.

- ***Promote shared responsibility across the AI value chain.*** Small software and device companies benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. The App Association urges NTIA's report and recommendations to reflect that all stakeholders developing and using AI have a shared responsibility for AI safety, efficacy, and transparency. AI policy frameworks, including those addressing dual-use foundation models, should ensure the appropriate distribution and mitigation of risk and liability (that those in the value chain who have the ability to minimize that risk based on their knowledge and ability to mitigate have appropriate incentives to do so).

  One way that NTIA could support shared responsibility is through proposing the creation of a mechanism for sharing best practices, and for surfacing timely threat indicators, similar to that employed by Information Security and Analysis Centers (ISACs), which foster information sharing across and between the government and private sector while avoiding liability for doing so.[6]

- ***Support, and rely on, international standards for risk management.*** The App Association supports reliance on international consensus standards to develop metrics for risk, creating standards for best practices, and/or supporting or restricting the availability of foundation model weights. We believe that NIST's approach taken in its AI Risk Management Framework is optimal. Support for and deference to international standardization would also align NTIA's efforts with the U.S. Government National Standards Strategy for Critical and Emerging Technology.[7]

- ***Coordination/Alignment with Other Leading Federal Efforts.*** Consistent with the intent of the Executive Order, alignment with other key federal efforts occurring in parallel should be prioritized. As a prime example, NTIA's recommendations should be consistent with the output of the U.S. AI Safety Institute.[8]

- ***Support international harmonization.*** We urge NTIA to maintain a priority for supporting risk-based approaches to AI governance in markets abroad and

[6] https://www.nationalisacs.org/about-isacs.

[7] https://www.nist.gov/standardsgov/usg-nss.

[8] https://www.nist.gov/artificial-intelligence/artificial-intelligence-safety-institute.

through bilateral and multilateral agreements. Already, developers of AI face top-down and one-size-fits-all mandates that substantially impede their ability to develop and utilize AI across a range of use cases. It is crucial that NTIA's efforts here, and the Administration's efforts broadly, discourage, or at least have a positive influence on, such mandates in other jurisdictions.

The App Association appreciates NTIA's consideration of the above (and appended) views and we urge NTIA to contact the undersigned with any questions or ways that we can assist moving forward.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

ACT | The App Association
1401 K St NW (Ste 501)
Washington, DC 20005
202-331-2130

Policy Recommendations for AI

ACT
The App Association

Artificial Intelligence (AI) is clearly a priority for policymakers, with 37 AI-related laws enacted globally, more than 80 pending legislative proposals at the state level and several more at the federal level. To understand and shape rules for this complex and evolving technology, a vital voice—that of small businesses, members of ACT| The App Association—must be prioritized in order to create a competitive, safe, and secure AI future.

We initially released these principles in 2021. However, we are updating them continually to reflect new developments in privacy and data security laws around the world and new learnings about the benefits, risks, and challenges presented by evolving AI tools in use cases from healthcare and education to software development and cybersecurity.

A successful policy approach to AI will align with the following guidelines:

1. Harmonizing and Coordinating Approaches to AI

A wide range of federal, local, and state laws prohibit harmful conduct regardless of whether the use of AI is involved. For example, the Federal Trade Commission (FTC) Act prohibits a wide range of unfair or deceptive acts or practices, and states also have versions of these prohibitions in their statute books. The use of AI does not shield companies from these prohibitions. However, federal and state agencies alike must approach the applicability of these laws in AI contexts thoughtfully and with great sensitivity to the novel or evolving risks AI systems present. Congress and other policymakers must first understand how existing frameworks apply to activities involving AI to avoid creating sweeping new authorities or agencies that awkwardly or inconsistently overlap with current policy frameworks.

2. Quality Assurance and Oversight

Policy frameworks should utilize risk-based approaches to ensure that the use of AI aligns with any relevant recognized standards of safety, efficacy, and equity. Small software and device companies benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. Policy frameworks addressing liability should ensure the appropriate distribution and mitigation of risk and liability. Specifically, those in the value chain with the ability to minimize risks based on their knowledge and ability to mitigate should have appropriate incentives to do so. Some recommended areas of focus include:

- Ensuring AI is safe, efficacious, and equitable.
- Encouraging AI developers to consistently utilize rigorous procedures and enabling them to document their methods and results.
- Encouraging those developing, offering, or testing AI systems intended for consumer use to provide truthful and easy-to-understand representations regarding intended use and risks that would be reasonably understood by those intended, as well as expected, to use the AI solution.

3. Thoughtful Design

Policy frameworks should encourage design of AI systems that are informed by real-world workflows, human-centered design and usability principles, and end-user needs. AI systems should facilitate a transition to changes in the delivery of goods and services that benefit consumers and businesses. The design, development, and success of AI should leverage collaboration and dialogue among users, AI technology developers, and other stakeholders to have all perspectives reflected in AI solutions.

4. Access and Affordability

Policy frameworks should enable products and services that involve AI systems to be accessible and affordable. Significant resources may be required to scale systems. Policymakers should also ensure that developers can build accessibility features into their AI-driven offerings and avoid policies that limit their accessibility options.

5. Research and Transparency

Policy frameworks should support and facilitate research and development of AI by prioritizing and providing sufficient funding while also maximizing innovators' and researchers' ability to collect and process data from a wide range of sources. Research on the costs and benefits of transparency in AI should also be a priority and involve collaboration among all affected stakeholders to develop a better understanding of how and under which circumstances transparency mandates would help address risks arising from the use of AI systems.

6. Modernized Privacy and Security Frameworks

The many new AI-driven uses for data, including sensitive personal information, raise privacy questions. They also offer the potential for more powerful and granular privacy controls for consumers. Accordingly, any policy framework should address the topics of privacy, consent, and modern technological capabilities as a part of the policy development process. Policy frameworks must be scalable and assure that an individual's data is properly protected, while also allowing the flow of information and responsible evolution of AI. A balanced framework should avoid undue barriers to data processing and collection while imposing reasonable data minimization, consent, and consumer rights frameworks.

7. Bias

The bias inherent in all data, as well as errors, will remain one of the more pressing issues with AI systems that utilize machine learning techniques in particular. Regulatory agencies should examine data provenance and bias issues present in the development and uses of AI solutions to ensure that bias in datasets does not result in harm to users or consumers of products or services involving AI, including through unlawful discrimination.

8. Ethics

The success of AI depends on ethical use. A policy framework must promote many of the existing and emerging ethical norms for broader adherence by AI technologists, innovators, computer scientists, and those who use such systems. Relevant ethical considerations include:

- Applying ethics to each phase of an AI system's life, from design to development to use.
- Maintaining consistency with international conventions on human rights.
- Prioritizing inclusivity such that AI solutions benefit consumers and are developed using data from across socioeconomic, age, gender, geographic origin, and other groupings.
- Reflect that AI tools may reveal extremely sensitive and private information about a user and ensure that laws require the protection of such information.

9. Education

Policy frameworks should support education for the advancement of AI, promote examples that demonstrate the success of AI, and encourage stakeholder engagements to keep frameworks responsive to emerging opportunities and challenges.

- Consumers should be educated as to the use of AI in the service(s) they are using.
- Academic education should include curriculum that will advance the understanding of and ability to use AI solutions.

10. Intellectual Property

The protection of intellectual property (IP) rights is critical to the evolution of AI. In developing approaches and frameworks for AI governance, policymakers should ensure that compliance measures and requirements do not undercut IP or trade secrets.