moz://a

Mozilla's Response to the Department of Energy's Request for Information on Frontiers in AI for Science, Security, and Technology (FASST) Initiative

About moz://a

Mozilla's mission is to ensure the internet is a global public resource, open and accessible to all. An internet that truly puts people first, where individuals can shape their own experience and are empowered, safe, and independent.

Founded as a community open source project in 1998, Mozilla consists of several organizations. Most notable is the non-profit Mozilla Foundation, which leads our movement-building work, and its wholly owned subsidiary, the Mozilla Corporation, which leads our market-based work, including the development of the Firefox web browser. They work in close concert with each other and a global community of tens of thousands of volunteers under the single banner: Mozilla.

For the past five years, Mozilla has been committed to advancing trustworthy AI. Mozilla published a paper in early 2024, Accelerating Progress Toward Trustworthy AI, that outlines how Mozilla and its allies are advancing openness, competition, and accountability in AI. Mozilla is putting its resources behind these priorities as well: The Mozilla Foundation has been dedicating 100% of its $30M a year budget to philanthropic activities, advocacy, and programmatic work on this topic. Mozilla is also investing another $30M in research and development on trustworthy AI via Mozilla.ai, as well as $35M in responsible tech startups — including startups with a focus on trustworthy AI — through Mozilla Ventures and the Mozilla Builders accelerator program. On the frontlines of modern AI practices, Mozilla freely provides an open-source, large-language model (LLM) AI model deployment system for local use and empowers more people to enhance the safety of models online through the 0Din bug bounty program. In addition, Mozilla supports the work of academic and civil society organizations engaged in artificial intelligence research, including the Open Source Initiative's work on defining Open Source AI, and through the work of the Columbia Convening which brought together a diverse set of stakeholders to talk about Openness and AI.

As an independent and mission-driven organization, Mozilla is committed to working with regulators to develop effective policies that ensure that innovation and growth in AI serve the public interest. In the past, Mozilla has provided comments on critical topics like NTIA's consultation on openness in AI models, NIST's AI Risk Management Framework, and NIST's Request for Comments related to Managing Misuse Risk for Dual-Use Foundation Models.

Executive Summary

Mozilla has been on the frontline of defending the open internet for 25 years and is a leader in the emerging conversation related to Public AI. On September 30, Mozilla released a paper outlining our view on Public AI, which elucidates Mozilla's vision for a robust ecosystem of initiatives that promote public investment, public orientation, and public interest AI throughout every step of AI development and deployment. Soon thereafter, Mozilla, along with the Economic Security Project and the Vanderbilt Policy Accelerator, hosted a convening to discuss the future of Public AI.

We applaud the Department of Energy's (DOE's) work in advancing the Frontiers in AI for Science, Security, and Technology (FASST) initiative. Such programs have the potential to create the foundation for Public AI infrastructure, which will not only help to enable increased access to critical technologies within the government that can be leveraged to create more efficient and useful services, but also have the potential to catalyze non-governmental innovation through the provision of resources such as public data sets. Along with programs like the National AI Research Resource, FASST could prove foundational to the nascent Public AI movement, ensuring that AI benefits all Americans.

Given DOE's leading role in technological innovation within the government through the national labs, NNSA, and advanced computing, DOE is well positioned for leadership in Public AI and FASST, helping to advance scientific discovery. In order for the DOE to maximize the potential beneficial impact of the FASST program and further catalyze private sector and civil society participation in the broader AI ecosystem, we recommend that the Department makes open source a key priority throughout the initiative from the creation of openly accessible unclassified scientific data sets to conducting red-teaming to detect threats and vulnerabilities with the support of the open source community. In addition to this high level recommendation, Mozilla has addressed many of the questions in the Request for Information below.

The following list highlights key themes from the questions addressed:

- **Benefits of Open Source:** Given Mozilla's long standing support of the open source community, a clear throughline in Mozilla's responses to DOE's questions is the importance of open source in advancing key government objectives. Below are four key themes related to the benefits of open source that appear in our responses:
    - **Economic Security:** Open source by its nature enables the more rapid proliferation of a technology and according to NTIA's report on Dual-Use Foundation Models with Widely Available Model Weights, "They diversify and expand the array of actors, including less resourced actors, that participate in AI research and development." For the United States, whose competitive advantage in global competition is its innovative private sector, the rapid proliferation of newly accessible technologies mean that new businesses can be created on the back of a new technology, speeding innovation, and that existing businesses, whether a

hospital or a factory, can more easily adopt new technologies, helping to increase efficiency.

  - **Expanding the Market for AI:** While costs are rapidly decreasing, the use of cutting edge AI products purchased from major labs and big tech companies are not cheap, and many small businesses, research institutions, and nonprofits would be unable to benefit from the AI boom if they did not have the option to use freely available open source AI models. This means that more people around the world get access to American built open source technologies, furthering the use of American technology tools and standards, forging deeper economic and technological ties, and ensuring that research institutions doing critical work in fields like drug discovery and cancer detection can leverage efficient and effective AI tools today.
  - **Security & Safety:** Open source has had demonstrable security and safety benefits. Rather than a model of "security through obscurity," open source AI thrives from having many eyes examining code bases and models for exploits by harnessing the wisdom of the crowd to find issues, whether related to discriminatory outputs from LLMs or security vulnerabilities.
  - **Resource Optimization:** Open source in AI means more than freely downloadable model weights – it means considering how to make the entire AI stack more open and transparent, from the energy cost of training to data on the resources used to develop the chips necessary to train and operate AI models. By making more information on AI's resource usage open and transparent, we can collectively work to optimize the efficiency of AI, ensuring that the benefits truly outweigh the costs.

- **Keep Competition Top of Mind:** The U.S. government wields outsized influence in shaping markets as its role not just as a promulgator of standards and regulations but due to its purchasing power. We urge the DOE to consider broader competitive concerns when determining potential vendors and partnerships for products and services ranging from cloud resources to semiconductors in an effort to foster a more competitive AI ecosystem, as emphasized in OMB's guidance to Advance the Responsible Acquisition of AI in Government which highlights the importance of promoting competition in procurement of AI. This will have far-reaching implications on the U.S. AI ecosystem, and building a more competitive industry is critical to making sure that America remains a global leader. The DOE should make an effort to work with a range of partners and civil society organizations rather than defaulting to standard government partners and big tech companies.
- **Making FASST "Public" By Default**: Mozilla has advanced a vision of Public AI, an ecosystem of initiatives that promote public investment, public orientation, and public use throughout AI development and deployment. It is critical that as FASST engages in the development of new models, datasets, and other tools and resources, it defaults to making its work public by default. This may mean directly open sourcing datasets and models, or working with partners, civil society, academia, and beyond to advance access to AI assets which can provide public value.

Questions & Responses

What kinds of data governance practices, risks, and opportunities should DOE take into consideration, particularly for open sourcing scientific corpuses to the community or interested parties?

It is important that DOE consider the importance of openness across the entire stack as it engages in data governance practices. In early 2024, Mozilla held a convening in partnership with Columbia University which brought together dozens of AI experts to discuss the intersection of openness and AI. From that convening and the proceedings documented in "Towards a Framework for Openness in Foundation Models," it was clear that AI should be looked at holistically and openness across the AI model stack including technical and non-technical artifacts, from model weights to code to model cards are critical for enabling good governance practices that benefit as many stakeholders as possible.

What are additional data-related tools and technologies DOE should invest in to promote AI-ready data and fuel continued US leadership in AI?

For the United States to continue to lead in AI, government support in the development of robust open access datasets will be extremely important given the criticality of data in enabling AI functionality. The DOE can further enable U.S. leadership by promoting the data collection and dissemination of datasets where the agency is uniquely positioned, such as the collection of scientific data from government labs, weather-related data, and materials science information. In order to disseminate the data effectively, DOE and DOE funded projects as well as DOE vendors should make their data open access and accessible for American researchers, civil society organizations, and enterprises.

How can DOE ensure FASST investments support a competitive hardware ecosystem and maintain American leadership in AI compute, including through DOE's existing AI and high-performance-computing testbeds?

In order for the FASST initiative to effectively support and ideally catalyze a more competitive AI hardware ecosystem, DOE should seek to purchase hardware from American startups and emerging hardware vendors whenever possible and serve as an early-adopter of new solutions. In addition, DOE should view the broader competitive tech and AI ecosystem when considering partnerships and purchases. This means that DOE should consider whether its actions may inadvertently support entities potentially engaged in anti-competitive behaviors throughout the AI stack.

How can DOE continue to support development of energy-efficient AI hardware, algorithms, and platforms?

Through the Mozilla Technology Fund, Mozilla has invested substantially in understanding the intersection of AI and energy use, funding research and open source tooling like deep learning energy measurement and optimization as well as Code Carbon, an open source software package which helps developers estimate and reduce the energy consumption and carbon emissions of any piece of code. In addition, Mozilla's Firefox is the only browser which has made it easy for users to understand expected carbon emissions based on usage via Firefox Profiler. We have found that one of the more difficult aspects of embarking on the research necessary to even understand AI and energy consumption is the lack of transparency, standardization, and access to relevant data.

As DOE through the FASST initiative develops new models and computing infrastructure, it would be immensely beneficial to the research community if DOE made data such as energy and resource consumption for AI training and inference public. It is critical that such data is made easily available, standardized, and for the entire AI lifecycle. This includes data from the energy intensity of semiconductor production to model training, to data points like daily water consumption for a data center, so that the research and developer community can help optimize energy and resource consumption. In addition, if DOE provisions compute and AI resources from private sector providers, requirements that providers make data on energy and resource use public would facilitate further research in the field and help to enable future discoveries and advancements in developing and deploying energy-efficient hardware, algorithms, and platforms. This would also help downstream service providers choose the most efficient AI system for their needs, potentially leading to both reduced costs and environmental impact. Finally, DOE can support the development of "small" models which are used for more specific tasks and are generally less energy intensive than training and inference for large language models.

How can DOE continue to support the development of AI hardware, algorithms, and platforms tailored for science and engineering applications in cases where the needs of those applications differ from the needs of commodity AI applications? How can DOE partner with other compute capability providers, including both on-premises and cloud solution providers, to support various hardware technologies and provide a portfolio of compute capabilities for its mission areas?

Having worked for many years as an open source company with millions of lines of code written in concert with a large open source community, we are keenly aware of the complexities and opportunities in gathering large groups to work towards a shared technological effort. The use of prizes present a unique opportunity to catalyze further innovation in spaces which may not have a direct or immediate commercial application, making it difficult for those developing in the space to

avail themselves of private sources of capital. Mozilla applauds the use of well-designed prizes and is keenly watching the myriad legislative initiatives including the AI Grand Challenges Act and Future of AI Innovation Act which seek to use prizes to accomplish similar objectives to DOE's. Mozilla suggests that DOE explore partnering with civil society organizations with expertise in running such contests, especially those with clean energy experience.

In addition, DOE should consider directly contracting civil society organizations and academic institutions with the necessary talent and capabilities in relevant areas to develop tailored science and engineering applications and consider open-sourcing tools created when appropriate to do so. This would help to augment DOE's capacity while having the added benefit of bolstering civil society and academic institutions and providing new tools that would allow other organizations and small businesses to better compete, creating a more dynamic and competitive AI industry.

How should DOE consider the benefits of open sourcing of scientific and applied energy AI models for the scientific community while fully addressing research security and other national-security concerns?

While the open sourcing of any model or resource can create risks, just as closed models can create risks, it is important to acknowledge the significant benefits of open source, especially when it comes to scientific and applied energy AI models. To start, open source allows for the rapid proliferation of a technology – a double edged sword – but one that plays to America's comparative advantage of economic dynamism and entrepreneurship by enabling a more competitive environment. According to a recent article in Just Security authored by Dr. Keegan McBride and Dean W. Ball, "...for many countries open source AI represents the only opportunity to engage with the technology due to its prohibitively high development and training costs." U.S. open source provides a low-cost alternative which helps to further American technical standards and commercial relationships.

How can DOE ensure foundation AI models are effectively developed to realize breakthrough applications, in partnership with industry, academia, and other agencies?

In order to make sure that the models developed by DOE are most effectively utilized in developing breakthrough applications, it is important that DOE approach the development and sharing of such models through the lens of Public AI. In practice, this could mean that DOE provides the relevant models open source, or under a license following best practices and guidelines set by the Open Source Initiative's definition of Open Source AI at a minimum, while ideally going further by providing training data as well. However, in order to make such breakthrough applications more accessible, DOE could, directly or through partnerships with civil society organizations, provide

technical assistance resources to support end users who wish to adopt and adapt models for breakthrough applications. For example, this could be a small academic research institution which may lack in-house AI development expertise but has a unique thesis and data set that are applicable for a potential breakthrough solution, who could collaborate with DOE and a civil society partner to effectively utilize a previously developed AI model for their needs.

DOE has an inventory of AI workforce training programs underway through our national labs. What other partnerships or convenings could DOE host or develop to support an AI ready scientific workforce in the United States?

DOE can better host an AI ready scientific workforce by more effectively collaborating with the U.S. open source community, which is generally technical and interested in developing technologies for the broader social good. In addition, DOE should consider collaborating with Regional Technology and Innovation Hubs, especially on topics related to the issue area the hub is meant to address, such as the Nevada Tech Hub which has a technology focus area of lithium batteries and electric vehicle materials. Finally, the Department of Energy should consider actively holding convenings in areas outside of Washington D.C., New York City, San Francisco, and Seattle, giving consideration to those areas that were previously commercial and scientific hubs in industries such as coal, metallurgy, and mining. DOE should also consider creating convenings and partnerships with academic institutions based outside existing tech hubs or other groups, such as unions that are actively seeking to create new talent bases to address critical technology needs, or academic institutions such as Dakota State University which is partnered with ArmyCyber and the NSA and Carnegie Mellon University which hosts the National Robotics Engineering Center.

How can DOE effectively engage and partner with industry and civil society? What are convenings, organizational structures, and engagement mechanisms that DOE should consider for FASST? & How can DOE improve awareness of existing allocation processes for DOE's AI-capable supercomputers and AI testbeds for smaller companies and newer research teams?

To improve the awareness of existing allocation processes for the Department's AI capabilities to smaller companies and research teams, DOE should partner with ecosystem connectors prevalent throughout startup, nonprofit, and research communities.

At Mozilla, we've seen the benefits of partnerships. We've supported the work of technology experts like Deb Raji and Ambassador Karen Kornbluh through the Mozilla Fellowship program, funded research on critical issues through the Mozilla Technology Fund, supported open-source AI startups through the Mozilla Builders program, invested in early-stage AI founders, and have engaged in public awareness campaigns like Mozilla's Open Letter to Meta regarding the shutdown of CrowdTangle, and much more. By collaborating with ecosystem connectors like Mozilla, venture capital firms, civil society organizations, and startups, DOE can more effectively reach a broader audience. DOE could also explore partnerships to embed access to allocation processes within other tools and services used by the AI community.

DOE can most effectively engage with industry and civil society by working proactively to incorporate such organizations into the agency's roadmap in a meaningful way. This could mean having a committee of leading AI companies, civil society organizations, and researchers, that work in tandem with DOE on FASST initiatives as core partners, it could take the shape of having industry and civil society commit resources to the FASST initiative following the example of the NAIRR, or a number of other engagement strategies. What is most critical is that outside partners are involved throughout the process as core stakeholders rather than as an afterthought. In terms of structure, the DOE could consider hosting bi-annual convenings of key stakeholders to discuss the progress and plans for FASST, creating a continuous forum for stakeholder input and accountability.

What role should public-private partnerships play in FASST? What problems or topics should be the focus of these partnerships?

When considering the role of public-private partnerships, the DOE should make sure to work with companies, academia, and civil society groups that are aligned with the overall mission of the organization, will further scientific progress, and which will actively develop products and solutions which will benefit American people. In addition, rather than focusing only on public-private partnerships with large AI labs and big tech companies, DOE should actively work with smaller startups and nonprofits in an effort to create a more competitive AI ecosystem whenever possible, but especially in critical areas such as semiconductor and hardware development and model development. As DOE considers potential partners, it is important that the Department considers the broader competitive dynamics which shape the current tech ecosystem in areas such as cloud computing, which may play an important role in the future stifling or enabling of a thriving domestic AI ecosystem.

Next, DOE should think about how it can use its role as a convener and funder as a catalyst for pushing the U.S. AI ecosystem forward, for example by working with stakeholder groups and partners to create more openly available data sets and tools. DOE should also look to involve outside organizations including those in civil society and academia as much as possible, including in red-teaming, organizing prize competitions, community management, dataset collection and

management, the organizing of convenings, and in the development of custom models. This will help to facilitate the development of a robust AI ecosystem necessary for American leadership not only in the development of leading AI solutions, but in global standards setting, education, and in creating resilient solutions.

Final Thoughts

Mozilla would like to applaud the Department of Energy for its work related to the Frontiers in Artificial Intelligence for Science, Security, and Technology (FASST) program. As DOE continues to develop the FASST initiative, it should keep in mind the clear benefits of open source, enhanced transparency, turning FASST outputs into public goods, and the importance of being cognizant as to the larger competitive environment in the AI ecosystem. Mozilla sees the FASST program as a step forward towards a vision of Public AI that benefits society writ large and is focused on delivering value for everyday people.

We hope that as DOE considers the answers to the RFI and further develops the program, it continues to consult and collaborate with a range of stakeholders, including in civil society, academia, and the open source community. Mozilla thanks DOE for the opportunity to provide our thoughts and is honored to be part of the process. We look forward to working with DOE and other stakeholders to ensure that the FASST initiative meets its broader objectives and creates significant benefits for all Americans.