PUBLIC SUBMISSION

Received: May 21, 2025
Tracking No. may-2l13-z8v1
Comments Due: May 28, 2025 **Submission Type:** Web

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0090
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Duality Technologies Inc.

General Comment

Attached please find Duality Technologies' comments regarding Docket ID No. NSF-2025-OGC0001. Thank you.

Attachments

OSTP NITRD AI RFI Docket ID NSF-2025-OGC0001 vSUBMITTED

Duality Technologies Response to Request for Information on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

Docket ID No. NSF-2025-OGC0001

Submitted by Duality Technologies Inc.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

Duality Technologies Inc. thanks the Office of Science and Technology Policy, the Networking and Information Technology Research and Development program, and the National Coordination Office for the opportunity to contribute the following comments to the above-referenced request for information.

Duality is at the forefront of Privacy Enhancing Technologies (PETs), empowering organizations to unlock the full potential of sensitive data through secure collaboration for advanced analytics and AI, while rigorously upholding security, privacy, civil liberties, and data sovereignty. Our platform offers a broad set of hardware and software privacy solutions to enable AI models to be safely build and deployed. The platform offers governance tools to manage how these privacy technologies are used – i.e., role and access controls, schema management, logging, and more – in order to ensure responsible use of AI, compliance, and auditability. This combination of capabilities, along with data science and analytical models, allow organizations to collaborate with one another on their most sensitive data in order to derive new insights while ensuring compliance with regulatory requirements and their own business policies.

Duality's founding team is comprised of world-renowned cryptographers, including Turing Award winner Prof. Shafi Goldwasser, Gödel Prize winner Prof. Vinod Vaikuntanathan, DARPA Fellow and co-founder of the HomomorphicEncryption.Org standards body Dr. Kurt Rohloff, and data science experts. We are focused on the challenge of secure AI, and have a deep history of working with the US Government. For example, we utilize OpenFHE, an open-source, standards-compliance Fully Homomorphic Encryption library built in part with US Government funding[1].

[1] https://www.openfhe.org/community/

We support the United States' Government's AI R&D goals in that they should help advance the country's position as the unrivaled and dominant world leader in the space, enhance economic and national security, and promote human flourishing. One of the key ways to advance these goals from an R&D perspective is to invest in technologies that will allow the United States to use the best quality data possible, wherever it may be located, and protect AI sufficiently so that it can be deployed safely on any data set.

Some experts estimate that high quality language and data sources used for training AI models could already be exhausted, or may be depleted in under 4 years[2]. This will leave quality data sitting in private data sets – owned by government agencies, corporations, and academia. At the same time, however, many AI efforts today in these sectors are hampered by security concerns. Government agencies and enterprises alike are restricting which prompts their users can deploy against generative AI providers, for fear of leaking sensitive information, and are hesitant to deploy AI and Generative AI against their more sensitive data due to privacy and security concerns. This will leave the United States in a challenging position of having high quality data on which to train and deploy AI, but not being able to access or use it. This blind spot will mean that the United States will not be able to build the best possible artificial intelligence capabilities nor make the best possible decisions.

Stated differently, solving these challenges would directly impact the administration's goals. Access to more diverse, high quality data sets will help the United States build an unassailably superior position in artificial intelligence. Those models that the United States will be uniquely positioned to build will be among the most effective, helping the country make the best possible decisions for economic purposes, national security goals, and for promoting human flourishing.

The goal of Privacy Enhancing Technologies is to solve this exact challenge, for different levels of security and privacy (e.g., those needed for an intelligence agency may differ from those needed for policing or public health). Prioritizing government R&D investment in these technologies is critical, because they will help the United States unlock the data it needs and build the best possible models in order to achieve its security and economic goals, and to support the flourishing of its citizens. Today, PETs are used in production in a variety of sectors, but more research is needed to support certain models and data scales - for example those used by LLMs, or for imagery analysis, and more. This is particularly true given more stringent security requirements – for example, there are some ways to protect privacy and security at scale when running an LLM, but not, for example, to the extent that it is post-quantum secure (i.e., resistant to attacks by nation state adversaries). This is particularly critical today, given geopolitical challenges and investments made by our

[2] https://www.nature.com/articles/d41586-024-03990-2

adversaries into quantum computers for the express purposes of breaking today's encryption. As such, an example of an R&D area of focus for national security would be to invest in the technologies and infrastructure that would make post-quantum secure LLM implementations scalable and performative.

Additionally, while some PETs are standardized, like Fully Homomorphic Encryption, others like Federated Learning are not, and no standards have yet been publicized by NIST. This too is a critical investment area, as standards would enable faster adoption in the public and private sector alike, and would enable different user groups to understand the risk tradeoffs when deploying secure and trustworthy AI, and help organizations overcome concerns of making their data available to AI models and vice versa.

Cars can only go as fast as they do because they have seatbelts and brakes. Likewise, investing strategically in the research, development, and standardization of Privacy Enhancing Technologies is not merely beneficial, but essential for the United States to harness the full power of AI. Just like cars need safeguards to enable us to get value from them, AI does too. As such, ensuring the secure, responsible, and widespread deployment of AI across both the public and private sectors, solidifying the United States' global leadership in this transformative technology, and harnessing the power of AI to promote the interests of the United States is only possible with the safeguards the PETs offer.