Army AI Layered Defense Framework Request for Information

Assistant Secretary of the Army for Acquisition, Logistics & Technology
Office of the Deputy Assistant Secretary of the Army for
Data, Engineering, and Software

1. Introduction

Purpose

This Request for Information (RFI) is issued by the Assistant Secretary of the Army for Acquisition, Logistics & Technology (ASA(ALT)) Office of the Deputy Assistant Secretary of the Army (DASA) for Data, Engineering, and Software (DES) to provide the United States (U.S.) Army with a better understanding of industry capabilities, potential sources, and best practices relevant to the definition and implementation of an Artificial Intelligence Layered Defense Framework (AI-LDF) for the U.S. Army. The AI-LDF is to be a thorough theoretical and practical framework for mitigating risks to AI Systems. The Army does not foresee establishing a single program to develop the AI-LDF. The goal is to improve the Army's methodology for building a comprehensive library of risks and mitigations unique to or inherent in AI systems which will inform and guide the development and implementation of subsequent AI models and software.

This RFI invites industry to submit relevant information, comments, capabilities, and recommendations for approaches, potential development, implementation opportunities, and corresponding business models.

Responses to this RFI will assist the U.S. Army in maturing the AI-LDF as part of an ongoing, collaborative dialogue between the Department of the Army, industry and academia and will continue to coevolve along with the field of artificial intelligence. After the receipt of responses to this RFI, the Government may invite some or all RFI respondents to further discuss their capabilities, solutions, and practices with Government representatives.

Intent

This is not a Request for Proposal (RFP), Request for Quotation (RFQ), or an invitation for bid, nor does its issuance obligate or restrict the Government to an eventual acquisition. All information received in response to this RFI will be used for market research purposes only. The

Government does not intend to award a contract because of responses to this RFI, nor otherwise reimburse respondents for the preparation of any information submitted or Government use of such information. Acknowledgement of receipt of responses will not be made, nor will respondents be notified regarding the outcome of the information received. Proprietary information in the submitted responses should be clearly marked.

By issuing this and future AI-LDF RFIs, the U.S. Army seeks to partner with industry to learn more about risks associated with traditional adversarial methods, such as Data Poisoning and Model Stealing, and emerging and future risks broadly associated all branches of computer science as well as the potential for security disruption from theoretical advances in future technologies such as quantum computing as well as developing and implementing industry-leading risk mitigation strategies and technologies.

2. Scope and Assumptions

The United States Assistant Secretary of the Army for Acquisition, Logistics, and Technology (ASA(ALT)) has issued this Request for Information (RFI) to enhance the U.S. Army's understanding of industry capabilities, potential sources, and best practices for implementing a multi-tiered Artificial Intelligence (AI) Layered Defense Framework. This framework aims to incrementally increase security measures from an open, accessible strategy to a highly secure approach with stringent controls, tailored to the sensitivity and importance of the data and models. The AI Layered Defense will serve as a thorough theoretical and practical framework for mitigating risks to AI models and software. Risk is broadly defined as the possibility that the occurrence of an event, related to AI, will adversely affect the achievement of the Army's objectives.

While AI software faces the traditional cybersecurity risks associated with all software, the Layered Defense Framework is concerned with building a comprehensive library of risks and mitigations unique to or inherent to AI: risk associated with the data used to train the model, the software/model itself, the use of the software, and the interaction of people, software and system. The AI Layered Defense Framework is intended to be a flexible, structured, and measurable approach to address AI risks prospectively and continuously throughout the AI lifecycle.

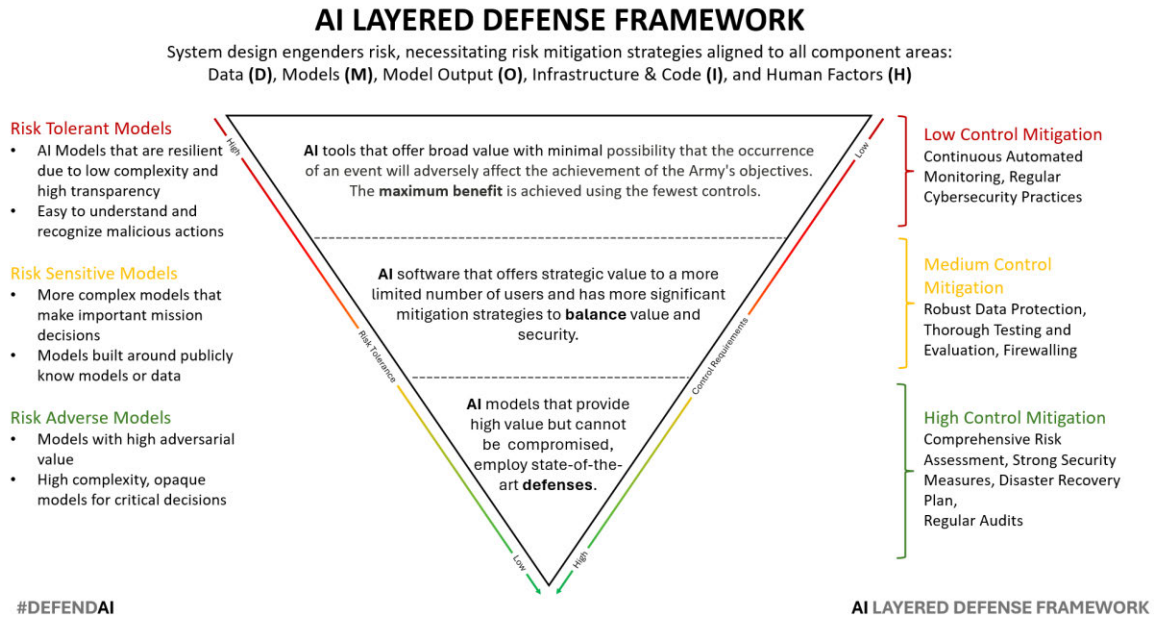

Figure 1. AI Layered Defense Framework

ASA(ALT) is interested in learning more about risks associated with traditional adversarial methods, such as Data Poisoning and Model Stealing, and emerging and future risks broadly associated all branches of computer science as well as the potential for security disruption from theoretical advances in future technologies such as quantum computing. Identifying risk is only the first step in developing and implementing industry-leading risk mitigation strategies and technologies. ASA(ALT) is committed to exploring computational methods for, among other things, detecting and removing "Trojaned" data among the vast public and crowdsourced data sets used to train models and detecting the creation of backdoors before deployment.

Additionally, the Army is open to exploring novel techniques for increasing model transparency and interpretability, providing a methodical approach to enhance the decision-making process, ensuring that outputs are not only efficient but also robust and justifiable. These are only two examples of the general areas of interest. ASA(ALT) is open to any proposed mitigation frameworks that will allow the Army to leverage the strengths of AI in handling large, complex and sensitive datasets and performing complex analyses while ensuring that the outputs adhere to the high standards required in defense operations.

This RFI contains the following components:

1. RFI introduction, scope and assumptions, response instructions, and disclaimer.
2. Artificial Intelligence Layered Defense Framework (AI-LDF) v1.0 (dated 30 July2024).

3. Response Instructions

Responses to this RFI are requested in the following two parts.

Response Part 1: Company Overview (limit one page)

Please provide the following company profile and demographic information:

1. Organizational Name, Address, and Country of Business (if organization has experienced name changes, please list all previous names used).
2. Company representative name, business title, and contact information.
3. Industry NAICS Codes (North American Industry Classification System) and business size for each NAICS code.
4. Relevant contract vehicles and contracts held (vehicle, agency, and expiration date) - please highlight U.S. Army and DoD contracts clearly.
5. Year company was established/founded.
6. Company ownership (public, private, joint venture). Provide information regarding Foreign Ownership, Control, or Influence (FOCI) issues associated with your company/and or offering(s).
7. Business Classification/Socio-Economic Status (e.g., large, small, 8(a), women-owned, hub-zone, Small Disadvantaged Business (SDB), Service-Disabled Veteran-owned).
8. Location of corporate headquarters.
9. Locations of facilities Outside of Continental United States (OCONUS), if applicable.
10. Location where currently incorporated.
11. Overview of products and services provided.
12. Overview of OCONUS experience with specific locations discussed, if applicable.

Response Part 2: AI-LDF Feedback (limit 12 pages)

Please provide open honest feedback on the initial AI-LDF as it pertains to its goal to develop a mutually-exclusive-and-collectively-exhaustive (MECE) list of AI specific risks and mitigations. The government would prefer to receive constructive criticism with recommendations for improving the AI-LDF.

In addition to any general feedback the Government requests responses to the following:

1. The AILDF incorporates elements of multiple risk frameworks (Databricks AI Security Framework, MITRE ATLAS, NIST AI RMF, IEEE AI RMF, Google's Secure AI Framework, Microsoft AI Risk Assessment for ML Engineers, IBM AI Lifecycle Governance, MIT Sloan Framework for assessing AI risk & others) –
    a. Are there other key sources of that should be considered?
    b. Are there critical risks that the framework does not account for?
    c. Are there mitigation strategies that aren't accounted for?
    d. Are there emerging theoretical or technological risk or mitigations that should be included?
2. The AILDF reflects an overlap between risk and mitigation strategies that are relevant to software/system development generally and those that are unique to AI enabled systems specifically.
    a. What strategies could be employed to identify, isolate, and address AI specific risks within the context of broader capability development efforts?
    b. Are traditional approaches to risk management sufficient for AI enabled system development? If not, what in your view are the key limitations? Does the AI-LDF in its current form adequately address these limitations? If not, how could it be improved?
3. How might the Army continuously monitor for, and incorporate, emergent AI risks?
4. Can you provide any examples of how you've implemented AI risk management into:
    a. Large scale development efforts?
    b. Business processes?
    c. Governance structures?
5. What strategies do you recommend for advancing the adoption of AI risk management practices?
6. What Red-Teaming strategies do you recommend for assessing the vulnerabilities of AI systems to adversarial attack?
7. Are there mitigations strategies that, when used individually or in concert with others, pose unique challenges in terms of technical feasibility?
8. Are there any considerations regarding AI, Risk, Mitigations, or strategy that we have overlooked when investigating adversarial threats to Army AI?

Response Submission

Responses should be submitted in Microsoft Word format or Adobe Portable Document Format (PDF). Responses should not exceed twelve (12) pages. A cover page is not included in the page count restriction. Responses should be in Arial font, 12 points. Respondents are strongly urged to adhere to this page limitation and limit marketing material to provide substantive information. Submit responses via e-mail to kye.h.park.civ@army.mil by 30, August, 2024 at 5:00 PM Eastern Standard Time.

4. Disclaimer

This RFI is issued solely for information and planning purposes. This RFI is not a solicitation and is not to be construed as a commitment by the Government to issue a solicitation or ultimately award a contract. Responses will not be considered as proposals, nor will any award be made as a result of this request. Federal Acquisition Regulation (FAR) clause 52.215-3, "Request for Information or Solicitation for Planning Purposes", is incorporated by reference. The Government does not intend to reimburse respondents for any costs associated with the submissions of their responses to this RFI; respondents to this RFI are solely responsible for all expenses associated with responding. Proprietary information and trade secrets, if any, must be clearly marked on all materials. All information received in response to this RFI that is marked "Proprietary" will be handled accordingly. Please be advised that all submissions become Government property and will not be returned nor will receipt be confirmed. In accordance with FAR 15.201(e), responses to this RFI are not offers and cannot be accepted by the Government to form a binding contract.

Responses from this RFI will be used to formatively shape broad U.S. Army guidance for acquiring software solutions. Response content may be aggregated and anonymously published into summary documentation to facilitate such guidance. Any publications resulting from this RFI will be non-attributional, and RFI respondents' consent for their responses to be used for such purposes.