CCC
Computing Community Consortium
Catalyst

The Computing Research Association (CRA)'s Computing Community Consortium (CCC) Response to the Defense Industrial Base Adoption of Artificial Intelligence for Defense Applications; Notice of Availability

Docket Number: DOD-2024-OS-0058-0001

This response is prepared by the Computing Research Association (CRA)'s Computing Community Consortium (CCC) by inviting CCC Council members and other members of the research community with interest and knowledge of the use of AI in defense applications to a roundtable discussion. The participants discussed the RFI and contributed to this written response document. CRA is an association of nearly 250 North American computing research organizations, both academic and industrial, and partners from six professional computing societies.

The mission of the CCC, a subcommittee of CRA, is to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges. Please note any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the authors' affiliations, or of the National Science Foundation, which funds the CCC.

July 22nd, 2024

Written by: Alex Aiken (Stanford University), Nadya Bliss (Arizona State University), David Danks (University of California, San Diego), Catherine Gill (Computing Community Consortium), David Jensen (University of Massachusetts Amherst), Daniel Lopresti (Lehigh University), Mary Lou Maher (Computing Community Consortium), Cristopher Moore (Santa Fe Institute), William Regli (University of Maryland), and Matthew Turk (Toyota Technological Institute at Chicago).

Infrastructure/Supply Chain Resilience

1. What foundational investments in the DIB does the DoD need to make to support increased adoption of AI into defense systems (e.g., manufacturing considerations, standards, best practices, bill of materials, etc.)? What foundational investments (e.g., standards, best practices, bills of materials, etc.) already exist within the DIB for defense systems that incorporate AI?

The DoD must establish teams of experts to conduct independent testing and evaluation of the models that DoD hopes to employ, to ensure these models are not at risk of exposing confidential, classified, or proprietary information. Significant investment should also be allocated to establishing facilities to conduct this independent testing, and to ensure these experts have the necessary resources to perform holistic and comprehensive evaluations.

The DIB also possesses significant engineering experience and risk management expertise that can be applied to the use of AI in technological settings. After these models have undergone significant independent testing and evaluation, the DIB should assign their existing experts to oversee the application of these models.

2. Are there specific vulnerabilities in the current and future supply chain that the DoD needs to address to support defense systems that incorporate AI?

Generative AI systems — including large language models and hybrid models — are developed using machine learning. *Such learning requires sets of training data, which should be considered a part of the supply chain, since they are a necessary component for creating these models*. Modern generative AI systems require such large amounts of training data that it is almost impossible to vet those data sets or to require that they come only from vetted sources. As a result, there is an opportunity for adversaries to introduce training data that could be used to trigger specific responses in such models. Such an approach could, in theory, create a “Manchurian candidate” scenario in which an apparently effective AI system can be triggered to produce specific responses when given specific input. Such scenarios have already been demonstrated in many systems based on smaller versions of the neural networks used to power modern generative AI systems. To date, such scenarios are almost entirely speculative, but they represent a serious potential vulnerability.

The AI “software stack”, as well as the people and organizations at the cutting edge of AI, are also not protected or functioning fully in the government's trust. The stack is brittle and complex, which creates system issues. These systems, and the people and organizations involved in developing them, must be fully secured and integrated into the DoD to ensure they are robustly protected from adversarial attacks and potential vulnerabilities.

Workforce

4. How can the DoD support the involvement of non-traditional defense contractors and small businesses in the design, development, testing, and deployment of AI technologies for defense applications?

Introducing AI into the DoD's standard operation will likely be very beneficial, especially for optimizing efficiency, but it also introduces significant risks. Because of this, comprehensive independent auditing from DoD classified auditors will be required, in addition to internal audits. Testing and evaluation of these AI models should not be avoided due to the highly classified nature of the data the models will operate on. Because this data must remain confidential it is necessary to put these AI models through substantial additional tests before they are implemented across the DoD.

The DoD should provide testbeds, test ranges, and shared facilities so that these capabilities do not need to be replicated separately by each outside contractor and small business. Providing these resources will significantly reduce the seed capital required to begin working on testing and implementing AI programs for defense applications. Shared facilities can also reduce the risk of data leakage by allowing data to be shared across LAN connections during testing and evaluation. For independent contractors that do not have access to secure DoD facilities, secure data access points must be established.

5. How can the DoD support and create effective partnerships with the DIB that will ensure that the DoD and DIB workforce is adequately trained, skilled, and sized to partner effectively?

AI and generative AI development is moving incredibly quickly, and keeping up with the speed of development is necessary for the DoD and DIB to consider. When the DoD began implementing Machine Learning, it encountered considerable issues with retraining personnel and overhauling systems with outdated or inefficient processes. As the DoD continues to implement AI, these issues will only become more prevalent. The DoD should rely on non-defense experts in the DIB, such as industry and academic leaders, who can advise the DoD on best practices to reskill and upskill current employees and train new employees. Non-defense experts, especially those who work on commercial AI models, can provide invaluable insight on advances and current trends in AI development, reducing the time it takes for the DoD to implement advanced techniques.

As for attracting new talent, the DoD and DIB should create greater career opportunities and pathways for advancement. Young professionals should be exposed to opportunities in AI development for public service early in their careers. The Federal Government can also support the DIB in creating a larger community of researchers on the forefront of this emerging field, similarly to the efforts it supported in developing a workforce for nuclear energy projects in the 1950s and aerospace projects in the 1960s.

Innovation

6. Are there specific intellectual property considerations or challenges related to the development of AI enabled defense systems that impact the DIB? If so, how can the DoD address these issues to promote innovation?

Because AI models are an emerging technology, regulation and legal considerations surrounding their use and ownership are evolving as well. However, in order to not stifle innovation, clear intellectual property standards must be established early on to allow non-defense contractors to engage on these projects with clear expectations. Open-source models are very popular among startups, including defense startups, but ownership questions still plague developers creating AI models which rely wholly or in part on open-source foundation models. For defense projects, which eventually may rely on AI models that include open-source code, it is necessary to ensure early on that the companies that own these open-source models will not call for these models to no longer be used for defense applications. Leaving the door open on this issue may result in damaging and complex legal battles in the future.

7. How can the DoD promote information-sharing and collaboration among government agencies, defense contractors, and research institutions to enhance data availability, collective knowledge, capabilities, and defense innovation in AI adoption into defense Systems?

The DoD should create shared resources, testbeds, and data lakes for use in developing AI applications. These resources should be created across classification levels, with individual users only having access to data up to their own classification level. Developing these resources will allow the DoD to not be dependent on industry companies for their data or data storage capabilities.

In addition, the DoD should promote translational research on federated and other types of distributed learning to have a better understanding of how simultaneous training from multiple entities affects an AI system. Further research should be conducted on methods for testing, data sharing, data harmonization, and utilization, to ensure that AI applications for defense are developed using optimal methods and techniques.

8. What measures can the DoD take to assess and mitigate the risks associated with potential adversarial exploitation of AI technologies within the DIB for developmental and/or operational defense systems?

Modern generative AI systems have unique issues with respect to adversarial exploitation, partially due to the opaque nature of these systems and partially due to their technical capabilities (the underlying technology of deep neural networks are universal function approximators, and thus can have extremely unpredictable performance in response to small changes in system input). AI systems also do not have a grounding in reality, so they are susceptible to misunderstandings or exploitations that seem obvious or nonsensical to human users and developers. To combat these shortcomings of AI systems, we suggest first that the DoD identify unacceptable and high-risk systems which should not be used, in order to limit the damage an adversary can cause by exploiting an AI system. Doing so also prevents relying on an AI system which, even without tampering by adversaries, may produce unreliable and potentially dangerous results.

AI systems that the DoD does deem acceptable for use in defense applications should also undergo regular and holistic testing, with special focus on identifying potential areas that can be exploited to introduce malicious data or may result in the leakage of classified information. New sandboxes must be created for testing these systems, especially when classified information is being stored or transported on non-defense servers (e.g., AWS). All of these risks must be considered as the AI systems are being developed, not as attacks occur. Adversarial attacks will always occur, and the best defense is to eliminate as many weaknesses in the systems as possible before they are deployed.

Acquisition, Policy, & Regulatory Environment

9. Please identify statutory, regulatory, or other policy barriers to the DIB's design, development, testing, and provision of AI-enabled defense systems in a manner consistent with DoD's approach to Responsible AI (https://rai.tradewindai.com/).

The largest issue with acquiring AI systems is that AI systems are never "finished". They are constantly evolving and require dedicated teams of experts to update and evaluate these systems. Because of the evolving nature of AI systems, acquisition models must also evolve and undergo evaluation regularly. IEEE is currently developing an AI procurement strategy to complement existing strategies created by the US Government Accountability Office, the World Economic Forum, and Ford Foundation. We strongly advise that the DoD and DIB rely on the IEEE AI acquisition guidelines, by referring to their IEEE GET Program and Standard for the Procurement of Artificial Intelligence and Automated Decision Systems. We also strongly recommend that the DoD and DIB refer to the AI Procurement Strategy, outlined by the Center for Inclusive Change, in determining contract considerations, transparency and explainability standards, and acceptable and unacceptable AI application areas and high risk systems.

12. What are the primary barriers that the DoD needs to address in the next five to ten years to enable the DIB to adopt AI for defense applications?

The DoD needs to provide clearer guidelines on its risk tolerance regarding AI systems and may need to adopt a slightly more risk-tolerant stance to ensure the DIB has a precise understanding of the requirements. Additionally, the DoD should offer indemnification for those developing AI applications for its use. Iit is also essential to embed DIB organizations more deeply within the DoD, to foster a more integrated understanding of the challenges. Furthermore, the DoD should elucidate what constitutes trustworthy and responsible AI in alignment with the directives they issue to achieve these goals.

13. In what ways can AI support or enhance acquisitions, supply chain management, regulatory compliance, and information-sharing in the DIB?

AI can be leveraged by the DIB to greatly improve the bureaucratic processes of the DoD. For example, AI can speed up finalizing and signing off on contracts by automating work and assigning responsibilities, such as reviewing and signing contracts, to those in charge, rather than requiring a human to oversee the process. AI can also automate information sharing, supply chain monitoring, note taking, reporting and summarizing, and other such necessary tasks that require a lot of time when handled by employees.