

July 22, 2024

Ms. Kylie Gaskins
Deputy Director
Enabling Future Capabilities Transition
Division of the Office of Policy, Analysis, and Transition
U.S. Department of Defense

Electronic Submission: www.regulations.gov, Docket No. 2024-11195

Re: NDIA Comments on Adoption of Artificial Intelligence for Defense Applications

Dear Deputy Director Gaskins:

The National Defense Industrial Association (NDIA) appreciates the opportunity to provide comments on the request for information on the Adoption of Artificial Intelligence for Defense Applications.

NDIA is the nation's oldest and largest defense industry association, representing over 1,700 corporate and over 65,500 individual members from small, medium, and large contractors, a majority of which are small businesses. NDIA members design, manufacture, apply, and maintain the cutting-edge technologies, systems, and platforms that our armed forces rely upon to deter aggression and defend our nation and its interests. As such, our members' professional and informed views on this proposed rule reflect the complexity and nuance of the issues under discussion.

Artificial Intelligence (AI) and Machine Learning (ML) are general-purpose technologies that can be leveraged across a broad range of use cases that offer tremendous benefits to our society and national security. Today, all major industries are utilizing AI to improve their product and service offerings to consumers, including everything from email spam filters to autonomous vehicles. Within the defense industrial base (DIB), industry innovators are partnering with the government to equip our warfighters with AI-enabled systems to improve the speed, quality, and accuracy of decisions in the field, which can provide the decisive advantage needed to deter or win a fight.

NDIA appreciates the Department's focus on AI, including the release of DoD's *2023 Data, Analytics, and Artificial Intelligence Adoption Strategy*. As the DoD continues to explore further advancements and uses of AI within defense applications, NDIA would emphasize the importance of collaboration with industry and other federal agencies and leveraging domestic and international consensus-based standard-setting organizations for AI to ensure harmonization and interoperability across the AI landscape.

NDIA would also offer an invitation to convene industry members to discuss the input provided below at the Department's convenience.

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 opportunities for applying AI technologies are effectively limitless. In order to help industry tailor its recommendations for AI investments that will be maximally effective in meeting mission needs, the Department should first focus on the following areas:

- **Defining AI:** Artificial intelligence is a broad category that ranges from relatively straightforward data analysis and search functions to highly complex threat prediction capabilities. Setting a clear, common definition for how the Department views AI as a technology enabler will allow businesses to provide clearer guidance for AI investments. The Department should leverage existing workstreams at the National Institute of Standards and Technology (NIST) to ensure harmonization across the AI landscape.
- **Setting clear direction on which mission areas are the highest priority for leveraging AI:** With such a broad range of possibilities for AI application, it is important for the Department to provide guidance on which mission areas it believes are most important and/or ready for AI investment, and which areas it either does not want or does not believe are essential for these efforts. Setting that direction will allow companies to apply their ingenuity toward solving our nation's most pressing challenges. The DoD should also incentivize the DIB to innovate with AI on existing programs, which require AI-driven processes and programs. When considering the acceptability of AI system outputs as deliverables, it is also essential to establish clear standards that the USG will accept under a contract.
- **Establishing a regulatory framework:** In order to move forward in suggesting and delivering AI solutions, companies will need to understand their compliance requirements. This includes but is not limited to considerations such as which AI is considered high-risk, data requirements, and overall governance requirements. Below are further considerations for potential elements of an AI regulatory framework:
 - NDIA supports a voluntary risk-based, use-case-specific approach and would again highlight the importance of collaboration with industry and other federal agencies on the development of any regulatory frameworks or approaches.
 - The DoD should also consider continuous research in the areas of training and testing AI that uses specific data sets.
 - The vendor community is often an AI developer, but the government is the AI deployer. Significant research needs to be done on the safeguards of AI algorithms during the development and training stages. We recommend that the Department clarify rules around use cases where a human in the loop (a responsible party) is required for high-

risk use cases at the deployment stage. Further, the government should create frameworks to ensure the AI models that are being developed have graceful degradation protocols if the model needs to be terminated.

- It will be important to note that existing technology-neutral laws and sector-specific regulations already apply to AI in the context of government procurement and cybersecurity policy.
 - Data privacy and security concerns should also be addressed. The DoD needs clear guidelines on handling sensitive data, especially with the potential for AI to be used with the kinds of information referenced in Executive Order 14017, "America's Supply Chains."
 - A regulatory framework needs to ensure that AI is fair and free of unintended bias. Fully unbiased AI is impossible to achieve. Rather, we should emphasize the need to test for and mitigate biases in data and algorithms.
 - A regulatory framework should also address explainability and interpretability to enable the AI ecosystem to better understand the functionality and trustworthiness of the AI system, including its outputs.
- **Use-Case-Specific Risk Frameworks:** Given the almost infinite range of potential use cases of AI systems across government, a generic, one-size-fits-all approach to regulation is impractical and could prevent or slow DoD's access to useful technology. Instead, policymakers should ensure that any proposal takes a risk-based approach that targets harms raised by specific applications of AI systems in high-risk use cases. Proposed regulations should focus on specifically defined use cases (rather than a general definition of "high risk") to enable clear legal analysis and an efficient development process.
 - **Independent Verification and Validation (IV&V):** DoD should consider investing in developing a framework for best practices in the IV&V of AI in defense systems. This could involve developing standardized approaches to evaluating AI risk in the system and developing playbooks of controls, CONOPS, architectural frameworks, and design patterns to mitigate certain types of risk associated with particular use cases. Risk levels can be managed, for example, through CONOPS or system architecture and design decisions, if these impacts are thoroughly understood. Examining approaches such as STPA (Systems-Theoretic Process Analysis), for example, for the use of AI in autonomous systems, can help in identifying potential hazards, understanding the interactions within the system, and establishing controls to prevent unsafe and undesired behaviors.
 - **DoD Evaluation:** The DoD could greatly help the DIB by continuing to identify U.S. industry partners from the burgeoning AI community willing to put proper protections in place to guard proprietary data and implement necessary security measures to foster the development of trained models and the sharing of those models within the DIB. Second, the DoD could consider approving and vetting existing AI companies and their existing tools for managing the supply

chain, implementing LLM, and applying generative multi-modal AI tools that could accelerate the development and safe operation of new capabilities for the DoD.

- **Publicly Available Data:** The DoD Chief Digital and AI Office (CDAO) recently communicated at the Advantage DoD 2024 AI symposium in March that CDAO will be opening access to government data over the next year, specifically by: a) publishing APIs to government data and federated data catalog, b) creating a hub for labeled data, and c) creating experimentation opportunities for industry get direct feedback on their products. The DoD should consider investments accelerate secure access to these datasets for DIB members through this initiative.
- **Foundation/Frontier Models:** The DoD should perform an assessment to evaluate the need to build its own foundation/frontier models and make these available to the DIB for incorporation into AI systems. While this may be a significant investment for the DoD, relying on commercial and limited-license open models creates risk and dependencies for DoD.
- **Toolkits and educational services for industry:** Existing foundational investments include the DoD CDAO's RAI Toolkit, which provides a centralized process that helps identify, track, and improve the alignment of AI projects with the AI Ethical Principles. The Toolkit offers tailorable and modular assessments, tools, and artifacts that span the entire AI product lifecycle, including phases such as intake, ideation, assessment, development/acquisition, testing and evaluation, integration and deployment, and use. The Toolkit includes resources and a list of tools to aid in responsible AI application development. The Toolkit is built upon established frameworks, including the NIST AI Risk Management Framework and the IEEE 7000 Standard, ensuring it incorporates industry best practices and ethical considerations.

DoD should also consider establishing an AI Expertise Center that could be a go-to resource for guidance, expertise, and best practices on AI geared towards start-ups and small and medium-sized businesses.

- **Training:** The DoD should invest in training programs to create a workforce with the skills to develop, deploy, and maintain AI systems. This includes data scientists, engineers, and ethicists.
- **Ensure access to commercial cloud resources:** DoD should explore ways to leverage existing discounting programs to ensure low cost for compute and inference services in future years. Computing power is essential for AI. Once the DIB integrates AI into its function and process, it cannot proceed without that level of computing. As airlines have done with fuel prices, DoD should consider investing upfront and locking in today's price for computing. Further, DoD should assess its access to commercial cloud computing capacity at all classification levels and ensure it has access to the requisite level of processing capacity at the unclassified, secret, and top-secret levels to handle mission-critical training and inference workloads.

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?

Meeting the needs of our national imperatives depends upon a diverse set of industry partners with access to critical technology. The Department should promote the use of commercial solutions rather than government-specific technologies and procure commercial technologies using commercial terms and conditions. Recent trends that threaten that balance include:

- **Access to chips and other technologies:** The ever-greater reliance on computing power to operate complex systems depends on access to chips and other key components. Large-scale commercial entities have the resources and buying power to monopolize this market if they choose or if, at any point, supplies become substantially limited. The Department should work to ensure that these technologies remain available to the defense industrial base to help meet mission needs as required.
- **Open-source and AI:** Today's landscape of AI capabilities increasingly includes open models and open datasets. DoD should increase its focus on open-source security and software supply chain risk management (SSCRM).

NDIA members find the lack of open models that enable free and permissive use concerning. All current major commercial open models contain significant license restrictions that explicitly deny the use of their models for defense applications, and we recommend DoD negotiate with the model creators to remove the "defense application" restriction. This is a major issue for the DIB and DoD equities. Most "open-source" models are open-weight models (e.g., the model weights are available, but the data and source code used for training these models are not visible, which makes it more difficult to investigate the integrity of the models themselves.

- **AI generated Code:** As coding itself is increasingly performed by AI, code written utilizing new AI coding assistants is susceptible to threat actors. This can involve placing "Easter eggs" that will lay hidden during inspection and only later execute when the AI coding tool starts creating code or otherwise running code in production. Unless the DIB and the DoD have carefully reviewed the source code of the 1.0 version and all subsequent patches of the AI coding tools prior to adding them onto secure systems and verified that they do not contain malicious instructions, a clever threat actor could bury within the AI tool's code any number of hidden instructions that might not be discovered before the damage is done. To be sure, this is a cybersecurity problem for all software being installed onto sensitive systems, but the complexity of AI coding tools may make these kinds of threats more difficult to find. NDIA recommends that the DoD ensure future AI policies take source code security and software supply chain security into account.
- **Hardware Supply Chain Security:** The future DoD computing supply chain is vulnerable due to the sheer amount of training compute resources. Given the private sector's rapid adoption of this technology, DoD may compete for scarce hardware resources in the future. To ensure the DoD's ability to deploy real-time AI capabilities, particularly Gen AI capabilities, we must increase investment in Gen AI computing capacity. Further, the DoD should also increase its

access to this critical hardware by leveraging commercial cloud capabilities at all classification levels. Additional investments in tactical (3U, 6U, etc.) compute resources beyond the GPU architecture may also be necessary to take advantage of the proliferation of Gen AI solutions across the DIB.

3. Are there specific sectors/subindustries within the DIB that face significant challenges in developing and applying AI to defense systems? If so, which sectors/subindustries are impacted and what challenges do the sectors/subindustries face?

- **Physical Goods:** Industries that have historically focused on physical goods production have not historically made significant investments in software, compute, and personnel that would enable them to apply AI to their work. However, these organizations and industries could realize significant benefits from the application of AI to optimize build processes, identify quality improvement opportunities, optimize maintenance schedules, and increase safety.
- **FedRAMP and the DoD CC SRG Process for Cloud-based AI:** The current FedRAMP authorization process is a bottleneck for AI, particularly with the fast-paced nature of AI development and the number of small businesses that are entering the space without the resources for full FedRAMP authorization. The Department should work closely with FedRAMP officials to help improve FedRAMP and AI deployment in the following ways:
 - Continue emphasizing the “approve once, use many times” ethos of the FedRAMP program to ensure ATO authorization package reuse across the DoD. This will result in efficiencies and cost savings for government buyers because they can focus on security as opposed to compliance. Many already-approved cloud services contain machine learning and AI capabilities, and the DoD should encourage their reuse across its enterprise.
 - To incentivize faster approvals of AI and GenAI solutions through the lengthy FedRAMP process, DoD should send funding to the General Services Administration to help offset their appropriated costs and increase approval throughput.
 - Work with NIST and the AI Safety Institute to support the development and publication of AI-specific security baselines that could make the assessment process much faster and easier.
 - Automate vulnerability scanning and compliance checks, which would free up time for the human experts to focus on the bigger picture.
 - Move to a continuous monitoring model to enable operators to keep AI systems up-to-date and secure without having to invest in expensive subject matter experts for re-authorization efforts.
 - Institute a flexible and risk-based approach to DoD AI authorization based on the sensitivity of the data involved.

- **Limiting Unintended Consequences:** All sectors/subindustries within the DIB face the challenge of limiting the authority and behavior of AI-enabled systems to avoid unintended outcomes and potential risks to the public while simultaneously reaping the potential benefits of AI-enabled DIB systems. Architectural best practices through government-furnished reference architectures for specific functions or AI-enabled use cases could help mitigate this challenge.

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?

- **Mentor-protégé programs:** The DoD should create mentor-protégé programs that allow existing contractors to bring expertise in meeting DoD customer expectations while leveraging small and innovative company capabilities. This may include specific SBIR and STTR funding targeted at nontraditional defense contractors. To be impactful at scale, such programs must allow mentor companies to have a number of protégés concurrently, and that number should take into consideration the size/scope of the mentor company. The current blanket limit of three protégés per mentor greatly reduces the potential impact of the program.
- **More SBIR, OTAs, and BAAs for AI:** The DoD should provide resources for nontraditional defense contractors and small businesses to understand how to work with the DoD and provide access to training data for AI developers. Smaller companies, particularly small and sophisticated software companies, may avoid partnering with the DoD because of the DFARS' antiquated data rights rules. These should be made more flexible for acquisitions of AI tools in order to attract these nontraditional defense contractors. NDIA recommends that the DoD continue encouraging nontraditional contracting methods.

The Department should promote the use of commercial solutions and procure commercial technologies using commercial terms and conditions. The DoD also needs to accept that commercial data rights and customer-generated IP rights are the future of faster innovation. Further, many nontraditional contractors are small and may not have money to engage in lengthy IP negotiations with the DoD. They have spent the money to develop their T&Cs. Asking them to negotiate again with the USG is too much, and the companies are likely to walk away. If the DoD wants innovation, it needs to accept commercial T&Cs and commercial IP terms.

- **Partnerships:** The DoD should continue to focus on government/industry partnerships with nontraditional defense contractors to accelerate the development of transformational capabilities but should also involve traditional DIB contractors in those efforts. The combined innovation of nontraditional contractors with the seasoned experience of more traditional DIB contractors would likely result in more resilient, operationally fieldable innovations through AI/ML that meet the DoD's specific needs and expectations.

Enabling flexibility facility security clearance onramps: DoD should strive to minimize existing regulatory and policy blockers to enable the development, deployment, and scale of AI capabilities across the DoD enterprise. For example, allowing flexibility in contract mechanisms that enable small DIB as well as large DIB members without prolonged facilities clearances (ex IL6) and connectivity authorization to applicable development environments for testing and validation are required for AI adoption. A repeatable, streamlined process that encourages both facility and Authority to Operate reciprocity (see FedRAMP answer above) will enable DoD to meet mission requirements.

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?

- **Consider new Consortiums and Public-Private Partnerships:** The establishment of DoD-industry consortiums will enable relationship development and collaboration on relevant issues. The Artificial Intelligence Safety Institute Consortium (AISIC) housed under NIST is a good example that unites AI creators and users, academics, government and industry researchers, and civil society organizations in support of the development and deployment of safe and trustworthy artificial intelligence (AI).
- **Training:** Joint training is beneficial in settings such as the Defense Acquisition University (DAU) with industry and DoD participants. The DoD should also conduct surveys within the DoD staff and DIB to understand the current state of AI awareness and capabilities and then build plans to close identified gaps. The Department should also increase training for both government and industry personnel to provide further education on the specifics of FedRAMP authorization for AI. Additionally, the Department needs to ensure that the acquisition community is trained to identify AI, develop AI evaluation criteria, and tailor contract provisions to address specific AI issues, such as specially negotiated license rights to address IP. Many small companies will choose not to engage in prolonged DoD negotiations or contracting because of the cumbersome process.
- **Staffing:** DIB finances do not support staffing top-tier AI developers in the quantity required. Foreign adversaries don't have this problem.
- **Security Clearances:** Provide security clearances for AI developers who work across multiple programs at various classification levels. As we work to stand up AI infrastructure capabilities to serve the DoD, we are often limited in our ability to provide services in classified domains due to limited cleared staff billets. Once the DoD tests a capability at the unclassified level, having more cleared personnel across the DIB will enable the DDO to move the solution to a classified fabric mission use.

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?

Developing AI-enabled defense systems often involves collaboration between multiple entities, including defense primes, nontraditional contractors, and research institutions. Determining ownership and sharing of intellectual property (IP) rights in collaborative projects is complex. We encourage DoD to leverage commercial terms and conditions whenever possible.

- **Protecting IP:** DoD must develop an IP policy that allows data sharing but protects IP. DoD's "take-all" approach to IP is not conducive to the non-defense primes—especially commercial companies/start-ups/nontraditional. Further, businesses should be protected from disclosing proprietary information, such as that which is subject to intellectual property or trade secret protection, in the context of AI.

With the incorporation of training data into AI/ML algorithms, it is currently unclear whether trained AI/ML models would be viewed as derivative works of their source data under U.S. copyright law. If seen as derivative works, a trained model must be trained on works with sufficient license permissions or face potential copyright infringement.

In addition, trained AI/ML models can irreversibly assimilate the source data they consume potentially transferring teachings of such data to the next customer or application. Unlike humans, a trained AI/ML model has a perfect memory of the last processed dataset and all those before it. For the DIB, this poses a new avenue of IP transfer that must be addressed to ensure that proprietary information from one company do not inadvertently migrate to another company by virtue of using the same AI/ML model. In some cases, the software industry's approach of anonymization and amalgamation with other data are not sufficient because of the proprietary teachings are inherently present in the data itself. These types of situations must also be considered when the U.S. government seeks to obtain broad rights in AI/ML models developed with U.S. government funds that might inherently contain a company's proprietary source data that was developed at private expense. In such cases, the government should acknowledge the data rights character of the embedded dataset in addition to the data rights character of the AI/ML model.

- **Determining Ownership:** Regulations must determine who owns the resulting models in a way that acknowledges the value of both the government-supplied data and the private entity's investment with associated model ownership. Regulations regarding ownership of the resulting models could follow the existing efficient model of intellectual property ownership residing within the industry. License rights and usage rights of these models should be structured to acknowledge both the value of the government-supplied data and the private entity's investment.

When dealing with hybridized data sets that combine USG data with other sources, it is crucial to develop a framework that clearly defines ownership and usage rights at the time that the contracting action is executed. This framework should ensure that both the government and private entities can benefit from the hybridized data sets while maintaining the integrity and security of the data.

- **Additional specific areas within IP that DoD should address include the following:**
 - DoD should implement a waiver of copyright liability risk for training data, especially for government data furnished to train a model. This waiver could utilize the copyright authorization and consent provisions of 28 U.S.C. 1498(b) and be modeled after the patent authorization and consent clause in FAR 52.227-1.
 - DoD should address the effect of using restricted contractor data as training data. For example, who would be able to utilize a model created using Limited Rights training data from multiple contractors?
 - DoD should address the warranty of accuracy of the output of the model. For example, who assumes liability if an AI generates a repair process that is faulty and, when followed, damages a product or causes injury to personnel?
 - RFPs and contracts should clearly define rights and responsibilities for training data, including stating which entity is responsible for the training data (provision, maintenance, accuracy). The DoD should allow the DIB to add disclaimers and limitations relating to training data, provide training data to the DIB for certain applications, and allow reasonable indemnities and limitations of liability related to training data.
 - Beyond model and data rights, there are opportunities for novel approaches and architectures being leveraged for the training, grounding, or evaluating of AI models. These novel architectures should be considered for their contribution to the broader AI capability.

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?

- **Consortiums:** As highlighted in question 5, the DoD could promote sharing and collaboration by creating consortiums with government, academic, and DIB participation where lessons learned could be shared in a closed environment, and specific DoD-related needs/challenges could be addressed.
- **Reform security protocols:** We need to reform security protocols to enable multi-layered security training environments. Currently, those few areas which can support multiple classifications/SAPs, are not also configured to support AI training.
- **DoD-generated data:** The DoD should establish a data strategy that leads to the establishment of datasets that may support an even playing ground for all aspects of the DIB. This will lead to novel AI architectures, as well as development and deployment processes and techniques. All data generated by DoD systems—in development and deployment—should be stored, mined, and made available for AI & ML. To allow for greater innovation, make all this data available to

the DIB. Include requirements for maintaining the history, provenance, and pedigree of data sets and models, as well as maintaining data/model traceability. Data product owners should be assigned by the DoD for all data sources. The data should be made centrally discoverable and accessible via APIs.

- **Additional ways to promote information-sharing and collaboration:**
 - Ensure that information is protected within the USG and from public disclosure.
 - Consider implementing an information-sharing safe harbor similar to that deployed for cybersecurity.
 - An example of when this worked was through the DoD's ongoing efforts to implement Zero Trust. The DoD provided funding to drive specified outcomes, tested processes and reporting, and industry was able to deliver capabilities. NDIA recommends that the DoD prioritize funding that promotes pilots that drive information-sharing and collaboration among all entities, which will identify aspects that work today and opportunities for improvement tomorrow.

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?

- **Industry Working Group:** AI technologies are many and varied in their implementation techniques and their governance considerations. The DoD should establish an industry working group to define the AI lifecycle components and interfaces of concern. This provides the DIB insights into the DoD concerns and enables industry to establish best practices for AI operations, security, and development as it pertains to defense systems applications.
- **Framework and standardized approaches:** As described under Question #2, develop a risk-based framework for the IV&V of AI in defense systems. Develop standardized approaches to evaluating AI risk in defense systems, including mitigating against adversarial exploitation of AI in operations. DoD can mitigate risks associated with adversarial targeting and exploitation of DIB AI technology applications by further expanding ongoing efforts to increase cybersecurity and protect DIB entities against adversarial AI-enabled cybersecurity threats. To enhance industrial cybersecurity, DoD should encourage DIB participants to aggressively leverage AI-powered cyber defense capabilities, such as those provided through the NSA Cyber Collaboration Center.
- **Approaches to mitigate risk:** Research and experimentation programs should place a primary focus on approaches to mitigate AI risks, including:
 - AI developers can share the testing and evaluation processes used to assess the safety and security of the system and how potential biases were tested for and mitigated.

- Data quality techniques to assess if training data sufficiently represent real-world distributions.
- Run Time Assurance (RTA) approaches.
- Formal methods and other approaches to prove the correctness of AI models.
- Enhance trust in AI systems through explainability and other techniques.
- Consider using DARPA’s Guaranteeing AI Robustness Against Deception (GARD) program’s Holistic Evaluation of Adversarial Defenses repository, which is available for interested DIB partners and researchers to take advantage of these resources
- Adversarial exploitation of AI technologies is directly related to the security and validity of the data used to develop the AI-enabled function and the system architecture in which it is implemented. Static systems that are trained and deployed will likely be the initial focus for mitigating potential adversarial exploitation since dynamic learning and reconfiguration in operation can lead to greater exposure to risk. DoD should focus on the human-machine teaming aspects of AI-enabled systems so the potential exploitation is mitigated by having a human “over the loop” or governing the final behavior of an AI-enabled system for the foreseeable future to ensure compliance with global humanitarian law governing the use of AI in a weapon system.
- Adversarial exploitation of AI-enabled systems has a lot in common with the “social engineering” of humans. Incorporate this view, in addition to normal risk mitigation approaches for software.
- Important to note that some of the leading computer research has shown that AI is currently more applicable and helpful to defenders than offenders. That means we should be using it. AI will increase the speed of finding vulnerabilities for both red and blue teams, thus operationalizing and detecting/patching, respectively. Creating self-healing systems that self-detect and mitigate operations that are out of specification will help the defenders. Random error checking, explainability (how an algorithm produces a certain result), and new advances in detection, when speech or code has been electronically generated, can help to detect anomalies.
- The DoD should take measurements to ensure that the risks associated with potential adversarial exploitation (for AI specifically) fall along the lines of effective training in equipping AI operators on how to identify and mitigate risks, i.e., prompt injection, data poisoning, etc. The DoD must equip the DIB with the right training to ensure that AI operators develop AI systems in a Secure, Reliable, Human Centered, Robust, and Equitable way and are trained in risk/exploitation identification.

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/>).

- DoD should continue to engage with industry in any efforts to develop and deploy AI policies or regulations.
- DoD should work to learn and apply nontraditional acquisition strategies. DoD acquisition should incentivize such strategies.
- Software acquisition is challenging given the length of the acquisition process—cannot expect to acquire solutions (bespoke or bolt-on) rapidly. DoD should also consider cloud-based, extensible systems at all classification levels to enable software to scale if they are able to deliver mission impact. Services need to be thinking about their acquisition strategies and allow flexibility over time.
- DIB should be allowed to use commercial cloud provider commercial regions for access to hyper scaling compute power to outpace our adversaries. Commercial regions have the broadest optionality and most up-to-date capabilities for the DIB. Security, encryption, governance, auditing, and visibility are at a point now to support this for more use cases, including R&D, experimentation, testing, and evaluation. DIB partners should be able to access these resources and manage the accounts without needing a government contract. If the USG requires a USG contract to gain access, then the development and deployment of AI/ML systems will always be behind. The DIB needs access to develop the systems and models and then sell them to the DoD.
- Airworthiness regulations are not written to accommodate AI/ML. AI/ML-related airworthiness standards are being developed but are largely focused on the needs of the civil aviation community. The level of airworthiness rigor and methods for gaining safety assurance for civil applications may not be appropriate for DoD use cases. Further investment by the DoD for changes to MIL-HDBK-516 and identification of acceptable performance-based risk analysis and safety assurance methods for AI/ML could foster more rapid development and deployment of AI-enabled capabilities.
- As mentioned above, there is a lack of consistent direction and standards on AI governance. Standard regulations is important for speed. We should not have a patchwork of different regulations depending on the agency or use-case.
- The federal government should consider the impact of the state patchwork of AI legislation for AI procured for the federal government and whether preemption is needed.

10. Please identify examples of DoD programs, strategies, policies, or initiatives that have provided effective support to the DIB in transitioning AI for defense applications. What made these programs, strategies, policies, or initiatives successful?

- VISTA X-62A enabled the world's first live aerial combat maneuvering between a fighter testbed flown by AI agents and a crewed F-16. The X-62A ACE team established a precedent for safely and responsibly testing flight-critical AI systems via a bounded (runtime assurance) framework, which paves the way for future aerospace development and certification in the United States.
- Government-furnished commercial secure cloud environments, platforms, and tools provided for DoD and IC programs have proven to be an effective strategy for supporting the DIB in transitioning AI for defense applications. The DoD should continue ensuring that DIB members can access a secure cloud at all classification levels to provide innovative AI and ML solutions to warfighters worldwide. DIB partners should be able to access these resources to own and manage the accounts without needing a government contract. If the USG requires a USG contract to gain access, then that means the development/deployment of AI/ML systems will always be behind. The DIB needs access to develop (at their own expense) the systems/models and then sell them to the DoD.

11. What DoD financing and acquisition mechanisms can help facilitate or incentivize the DIB to continue to invest in AI technologies for defense applications?

- There is a focus on attrition-tolerant systems and rapid deployment of AI-enabled capabilities at low costs, such as the Replicator program. However, these smaller, lower-cost systems may not be capable of meeting the more demanding mission needs of larger platforms, and the AI-enabled capabilities they employ may not be scalable to other use cases or missions. Funding direct IRAD programs in partnerships with DIB contractors would allow investment in AI technologies for defense applications at lower individual risk and investment compared to individual company investment alone.
- AI technologies often require a platform approach to development, requiring significant investment. As such, it would be prudent to establish financing mechanisms that support not just the development of models and algorithms but also the development of software/AI systems needed to continuously enhance and develop new AI models and capabilities.
- DoD should unify expectations across the Services. As noted in response to question #1, data availability is important.
- As noted in response to question #6, intellectual property protections and associated funding are needed to ensure private entities will invest in developing new models and in training models for usage in DoD systems and services.
- DoD should continue supporting commercial cloud. This must include the government's ability to acquire multi-tenant capabilities provided by SaaS vendors, as well as commercially available

marketplace capabilities. This involves partnering with the Defense Pricing Contracting and Acquisition Policy to formalize policy, mirroring that of GSA,¹ to allow for multi-year procurement of cloud services.

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?

- At the moment, the greatest challenge to AI adoption is the lack of policy guidelines to inform industrial base approaches. A consistent, standard understanding of the outlines of what DoD believes are acceptable—or unacceptable—approaches to AI use will unlock creative solutions from across the industry. The Department should develop a clear set of “guardrails” for implementation and then adopt an approach that empowers companies to advance AI approaches so long as they are not prohibited.
- The DoD needs to lead an honest assessment of the technical maturity of AI-enabled capabilities, specifically for autonomy, data analytics, perception, decision-making, and other areas to help cut through the hype being generated by the broader AI industry. Lessons learned from the self-driving car industry and the promise of “fully autonomous” capabilities have shown there may need to be an interim step utilizing AI-enabled systems with human and machine teaming for better overall performance in lieu of pushing to replace the human. While that goal is potentially reachable for some specific use cases, AI may be over-promising and under-delivering, causing distractions in areas where incremental improvements could be transformative.
- As in response to question #11, AI technologies often require a platform approach to development. Without consensus /guidance on the necessary AI operations required to meet the "trust" threshold for DoD applications, R&D efforts in AI will be insufficient for the deployment and fielding of these capabilities.
- As noted in response to question #6, intellectual property protections are crucial to supporting needed private investment in model training and advancement.
- Identify pathways to leverage international ally and partner investments in AI.
- DIB does not have the money required to hire top-tier AI talent. This gap will eventually force the defense industry to rely on commercial derivative technology only, which may or may not be sufficient to defeat adversaries with top-tier talent.
- Defense data is often classified and siloed, hindering collaboration and AI training on robust datasets. The DoD needs to establish secure data-sharing mechanisms while upholding national security.

¹ https://www.gsa.gov/system/files/MV-21-06%20with%20sup%201_0.pdf

- The DoD needs to attract and retain a skilled AI workforce, including data scientists, engineers, and ethicists, to compete with the private sector.
- Integrating AI with existing defense infrastructure can be complex and expensive. The DoD needs to find ways to modernize legacy systems or develop new AI-compatible architectures.

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

- AI, specifically Gen AI, has demonstrated its ability to accelerate productivity in nearly all industries. Acquisitions will benefit from accelerated insights, while supply chain management will benefit from integrated data insights and analytics that provide efficiencies.
- AI has the potential to augment the entire acquisition cycle. AI-enabled proposal writing, cost estimating, schedule planning, supplier control, compliance planning, model development, and artifact generation for regulatory compliance are all areas where having DoD standards and/or example models in place that could be shared with the DIB could have a dramatic positive impact. The use of AI in this process could fundamentally change acquisition products and evaluation activities.
- The DIB should be informed if AI products were used in the creation of RFP materials or if they will be planned for use in the evaluation. The government should also disclose when it's using AI for source selection or another way in the market research or other proposal processes.
- AI can assess and identify correlations and insights in large data sets that can support more effective decision-making, identify compliance hot spots, and allow for a more efficient allocation of resources.
- The DoD should support innovative advancements while leveraging existing commercial technologies.
- The Department should focus on leveraging cloud-based solutions at all classification levels.
- Streamline the FedRAMP authorization process by leveraging tools that automate evaluation and risk identification. For example, AI can be used in tools that enable true continuous monitoring, reporting, and threat mitigation, thus enhancing cybersecurity across the enterprise.

Conclusion

NDIA and its membership appreciate the government's desire to promote a strong, dynamic, and robust defense industrial base. If you have any questions related to these comments, please contact Michael Seeds at mseeds@NDIA.org.

Sincerely,

National Defense Industrial Association