

The Department of Defense Posture for Artificial Intelligence

Assessment and Recommendations for Improvement

Once again, artificial intelligence (AI) is at the forefront of our collective imaginations, offering promises of what it can do to solve our most challenging problems. As the news headlines suggest, the U.S. Department of Defense (DoD) is no exception when it comes to falling under the AI spell. But is DoD ready to leverage AI technologies and take advantage of the potential associated with them, or does it need to take major steps to position itself to use those technologies effectively and safely and scale up their use? This is a question that Congress, in its 2019 National Defense Authorization Act (NDAA), and the Director of DoD's Joint Artificial Intelligence Center (JAIC) asked RAND Corporation researchers to help them answer. This research brief summarizes that report.

RECOMMENDATIONS

- DoD should adapt AI governance structures that align authorities and resources with its mission of scaling AI.
- The JAIC should develop a five-year strategic roadmap that is backed by baselines and metrics.
- Each centralized AI service organization should develop a five-year strategic roadmap also backed by baselines and metrics.
- There should be annual or biannual portfolio reviews of DoD investments in AI.
- The JAIC should organize an annual or biannual workshop that showcases DoD's AI programs.
- DoD should advance the science and practice of verification, validation, testing, and evaluation (VVT&E) of AI systems.
- All funded AI efforts should include a budget for AI VVT&E.
- All agencies within DoD should create or strengthen mechanisms for connecting AI researchers, technology developers, and operators.
- DoD should recognize data as critical resources.
- The chief data officer should make some DoD data sets available to the AI community.
- DoD should embrace permeability—and an appropriate level of openness—as a means of enhancing DoD's access to AI talent.

Artificial Intelligence and DoD

The term *artificial intelligence* was first coined in 1956 at a conference at Dartmouth College that showcased a program designed to mimic human thinking skills.¹ Almost instantaneously, the Defense Advanced Research Projects Agency (DARPA) (then known as the Advanced Research Projects Agency [ARPA]), the research arm of the military, initiated several lines of research aimed at applying AI principles to defense challenges (see Figure 1). Since the 1950s, AI—and its subdiscipline, machine learning (ML)²—has come to mean many different things to different people: For example, the 2019 NDAA cited as many as five definitions of AI, and no consensus emerged on a common definition from the dozens of interviews conducted by the RAND team for its report to Congress.

To remain as flexible as possible, the RAND study was not bound by precise definitions, asking instead, “How well is DoD positioned to build or acquire, test, transition, and sustain—at scale—a set of technologies broadly falling under the AI umbrella”? And if those technologies fall short, what would DoD need to do to get there?

The RAND team distilled the NDAA’s mandate into the following three guiding questions:

- What is the state of AI relevant to DoD?
- What is DoD’s current posture in AI?
- What internal actions, external engagements, and potential legislative or regulatory actions might enhance DoD’s posture in AI?

For the first question, the RAND team purposely avoided trying to determine what technologies DoD should pursue or how DoD currently measures up to other countries in terms of AI uptake because that was

outside their mandate. Rather, the team assessed what DoD decisionmakers need to know about AI.

For the second question, the team assessed the posture of DoD for AI using a framework of six dimensions (Table 1).

What Decisionmakers Need to Know About AI

Examining the implications of AI for DoD and strategic decisionmaking requires taking a holistic view that considers three critical elements and how they interact:

- the technology and capabilities space
- the spectrum of DoD AI applications
- the investment space and time horizon.

The technologies and capabilities space: This includes the approaches, such as algorithms, that underpin current AI solutions. Although many technologies underpin AI, current interest (and hype) is fueled by advances in a small number of areas, like deep learning.

But success in deep learning requires large data sets. Deep learning algorithms tend to be highly specific to the applications for which they were designed, and demonstrated applications have tended to be commercial. What’s more, VVT&E remains very challenging across the board for all AI applications, including safety-critical military applications.

The spectrum of DoD applications: The spectrum of DoD AI applications can be characterized by where those applications fall in terms of four independent factors: operating environment, resources, tempo, and implications of failure (see Figure 2). The position on this spectrum can be summarized in terms of three overlapping bins:

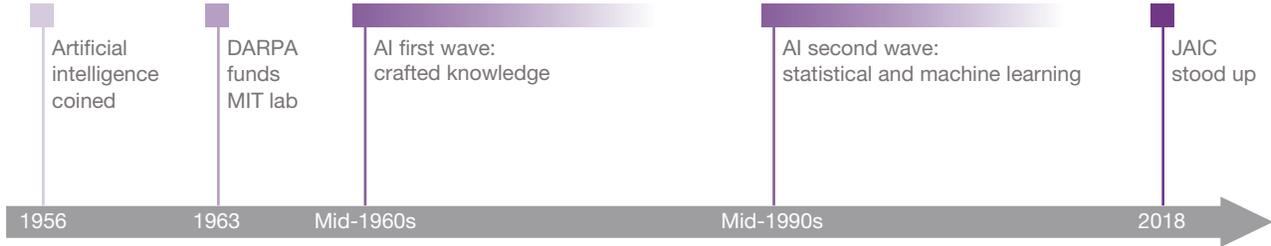
- *enterprise AI*, consisting of such applications as the management of health records at military hospitals in well-controlled, slower-paced environments, where analysts and decisionmakers have access to ample computational resources, data might be considered recoverable, and implications of failure might be negligible
- *mission-support AI*, consisting of applications like the Algorithmic Warfare Cross-Functional Team, also known as Project Maven, which aims to use ML to assist humans in analyzing large quantities of imagery from full-motion video data collected in the battle theater by drones

What Are the Recent Leap-Ahead Advances in AI and ML?

Deep learning, a form of supervised machine learning, trains many-layered neural networks on large sets of labeled *training data* to accomplish classification and prediction tasks, thereby offering new solutions to long-standing image, text, and speech-recognition problems.

Deep reinforcement learning uses deep learning to solve for approximate solutions of sequential decisionmaking problems with deferred rewards and has applications in strategic and computer games.

FIGURE 1
A Brief History of AI and DoD

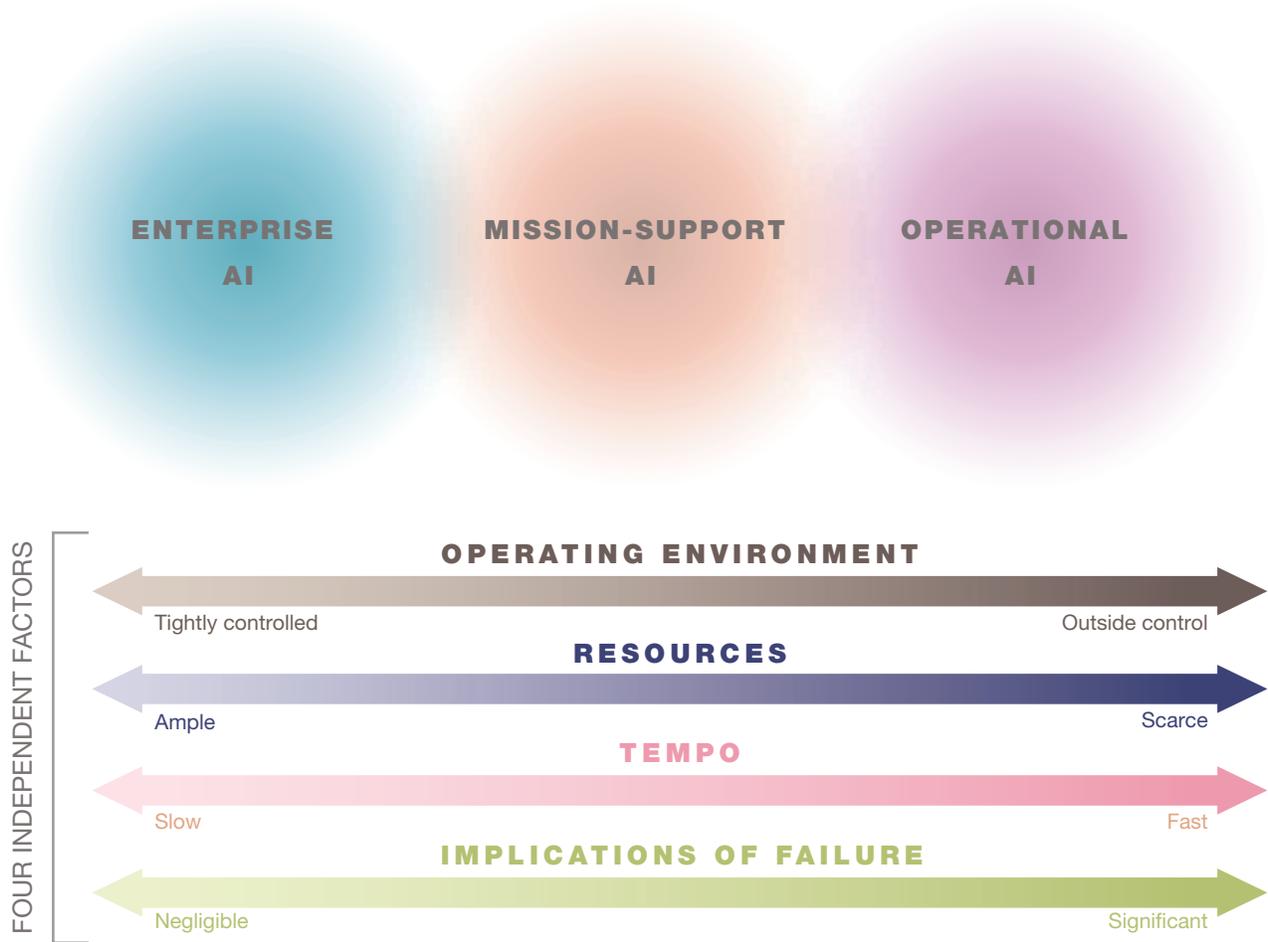


NOTE: MIT = Massachusetts Institute of Technology.

TABLE 1
DoD’s Posture for AI Was Assessed Across Six Dimensions

Dimension	Components
Organization	<ul style="list-style-type: none"> • Vision, strategy, and resource commitments • Organizational structures • Stakeholders and their mandates, authorities, and roles
Advancement	<ul style="list-style-type: none"> • Research and development portfolio and activities • Prototyping • VVT&E
Adoption	<ul style="list-style-type: none"> • Procurement • Fielding, sustainment, and life-cycle management • Development of doctrine; concepts of operations; tactics, techniques, and procedures; and processes
Innovation	<ul style="list-style-type: none"> • Internal culture of innovation • Mechanisms for leveraging external innovations • Mechanisms for engaging external innovators
Data	<ul style="list-style-type: none"> • Data as a resource • Governance of data collection and use • Storage, computing, and other infrastructure
Talent	<ul style="list-style-type: none"> • Talent needed to develop, acquire, sustain, and operate • Recruitment, retention, cultivation, and growth • Career management

FIGURE 2
Spectrum of DoD AI Applications



- *operational AI*, consisting of applications of AI integrated into weapon systems that must contend with dynamic, adversarial environments; fast tempo; scarce computational and communication resources (and possibly data); and significant implications of failure for casualties and risks to strategic objectives.

The investment space and time horizon:

In addition to the investments needed to develop or acquire AI technologies across the spectrum of applications, success in AI requires the following three other kinds of investments:

- *technological and other enablers*, such as infrastructure to enable the collection and management of data
- *VVT&E foundations and practice* for technological checks and balances
- *foundational basic research* that is not specifically aligned with a particular product or application to maintain longer-term technological superiority.

Finally, to manage expectations and ensure continued support, it is important to set realistic goals for the lead times that AI will need to progress from demonstrations of what is possible to full-scale implementations in the field. The RAND team’s analysis suggests that, as a rule of thumb, sustained DoD investments made as of the time of the original report’s publication (2019) can be expected to yield at-scale deployments in the

- *near term* (up to five years) for enterprise AI
- *middle term* (five to ten years) for most mission-support AI
- *far term* (longer than ten years) for most operational AI applications.

Of course, DoD can expect and should pursue faster progress for some applications, even within operational AI. However, these timelines reflect the RAND team’s assessment of what reasonable expectations are, given the current state of the technology and taking into account the four factors discussed previously: operating environment, resources, tempo, and implications of failure.

DoD’s Posture for AI

Overall, the RAND team found that, despite some positive signs, DoD’s posture is significantly challenged across all dimensions of the posture assessment.

- Organizationally, at the DoD level, the current DoD AI strategy lacks both baselines and metrics for assessing progress. Thus far, the JAIC has not been given the authority, resources, and visibility needed to scale AI and its impact DoD-wide. Similar challenges are seen at the level of the individual services.
- Data are often lacking, and when they exist, they often lack traceability, understandability, accessibility, and interoperability.
- The current state of VVT&E for AI technologies cannot ensure the performance and safety of AI systems, especially those that are safety-critical.
- DoD lacks clear mechanisms for growing, tracking, and cultivating AI talent, a challenge that is only going to grow with the increasingly tight competition with academia, the commercial world, and other kinds of workspaces for individuals with the needed skills and training.
- Communications channels among the builders and users of AI within DoD are sparse.
- Current DoD practices and processes—or their implementation—might be hampering innovation within DoD and inhibiting DoD’s ability to bring in external innovation.

Despite some positive signs, DoD is poorly positioned across all dimensions to effectively leverage and scale AI. The report offered a set of recommendations, both strategic and tactical, for moving forward.

Recommendations

The RAND team's recommendations to DoD are as follows:

- DoD should adapt AI governance structures that align authorities and resources with its mission of scaling AI.
- The JAIC should develop a five-year strategic roadmap that is backed by baselines and metrics.
- Each of the centralized AI service organizations should develop a five-year strategic roadmap that is backed by baselines and metrics.
- Annual or biannual portfolio reviews of DoD-wide investments in AI should be led by the JAIC, in partnership with the Under Secretary of Defense (USD) for Research and Engineering (R&E), the USD for Acquisition and Sustainment (A&S), the Chairman of the Joint Chiefs of Staff, and the service AI representatives.
- The JAIC should organize an annual or biannual technical workshop that showcases AI programs DoD-wide.
- DoD should advance the science and practice of VVT&E of AI systems, working in close partnership with industry and academia. The JAIC, working closely with the USD (R&E), the USD (A&S), and Operational Test and Evaluation, should take the lead in coordinating this effort, both internally and with external partners.
- All funded AI efforts should include a budget for AI VVT&E.
- All agencies within DoD should create or strengthen mechanisms for connecting AI researchers, technology developers, and operators.
- DoD should recognize data as critical resources, continue instituting practices for their collection and curation, and increase sharing while resolving issues in protecting the data after sharing and during analysis and use.
- To spur innovation and enhance external engagement with DoD, the chief data officer should make some DoD data sets available to the AI community.
- DoD should embrace permeability—and an appropriate level of openness—as a means of enhancing DoD's access to AI talent.

Notes

¹ Ipke Wachsmuth, “The Concept of Intelligence in AI,” in Holk Cruse, Jeffrey Dean, and Helge Ritter, eds., *Prerational Intelligence: Adaptive Behavior and Intelligent Systems Without Symbols and Logic*, Dordrecht, Netherlands: Springer Science and Business Media, 2000. This paper suggests that the program was actually developed at RAND.

² *Machine learning* designs algorithms to identify patterns in large data sets.

This brief describes work conducted within the Acquisition and Technology Policy Center of the RAND National Security Research Division (NSRD), which operates the National Defense Research Institute (NDRI), and documented in *The Department of Defense Posture for Artificial Intelligence: Assessment and Recommendations*, by Danielle C. Tarraf, William Shelton, Edward Parker, Brien Alkire, Diana Gehlhaus, Justin Grana, Alexis Levedahl, Jasmin Léveillé, Jared Mondschein, James Ryseff, Ali Wyne, Dan Elinoff, Edward Geist, Benjamin N. Harris, Eric Hui, Cedric Kenney, Sydne Newberry, Chandler Sachs, Peter Schirmer, Danielle Schlang, Victoria M. Smith, Abbie Tingstad, Padmaja Vedula, and Kristin Warren, RR-4229-OSD, 2019 (available at www.rand.org/t/RR4229). To view this brief online, visit www.rand.org/t/RB10145. The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. **RAND**[®] is a registered trademark.

Limited Print and Electronic Distribution Rights: This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents for commercial use. For information on reprint and linking permissions, please visit www.rand.org/pubs/permissions.

© 2021 RAND Corporation