GovCIO MEDIA & RESEARCH

DeepDives

DEFENSE AI INNOVATION

INSIDE:

SPONSORED BY

THUNDERCAT TECHNOLOGY

NetApp

GovCIO MEDIA & RESEARCH

From the editor's desk

Amy Kluber, Editor-in-Chief

Defense Needs AI

The Defense Department is all in on AI. Defense leaders have consistently touted AI's benefits in helping the department realize its vision for concepts like Joint All Domain Command and Control (JADC2), combatting future cyberspace wars and unlocking more bandwidth for personnel to execute mission-critical tasks while automating more routine ones.

But getting there is not without its challenges. While various pockets of innovation exist within the defense ecosystem, such as the Navy Research Lab's effort to put robots on ships for everyday routine maintenance and monitoring, officials are also focused on best practices for acquiring the technology and integrating it ethically.

This is where starting with the data will be critical. Industry leaders working on many use cases within this technology in the defense context see data computing at the edge to be a game changer for missions like predictive maintenance and connectivity.

Table of Contents

Amy Kluber,
Editror-in-Chief

Anastasia Obis,
Staff Writer/
Researcher

Navy Wants Robots to Perform Routine Shipboard Tasks

The Naval Research Laboratory has been building robots to perform ship maintenance and repair fragile and expensive satellites to help sailors prioritize mission needs.

BY ANASTASIA OBIS

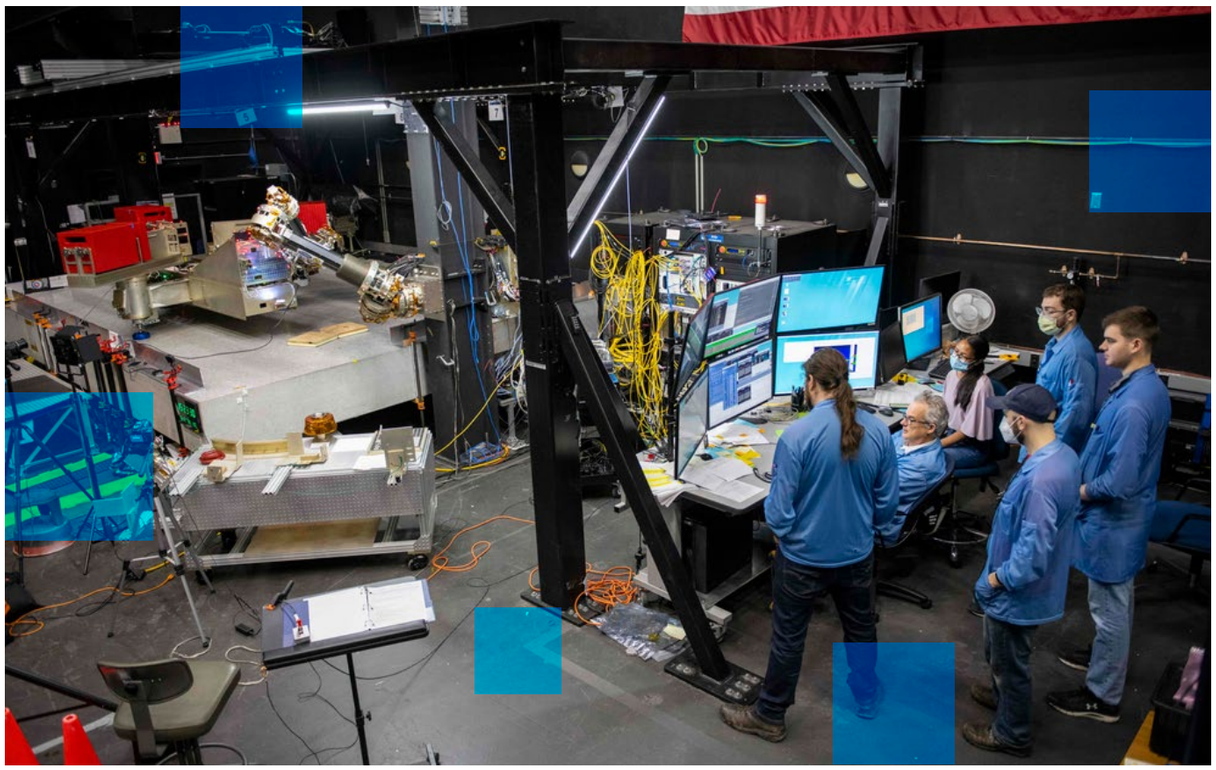

The U.S. Naval Research Laboratory's team analyzes data collected from contact dynamic testing on the robotic test bed in Washington, D.C. June 16, 2022.

Photo Credit: Sarah Peterson / DVIDS

The U.S. Navy may soon get robots to perform routine maintenance and janitorial tasks on Navy ships — alleviating the burden of monotonous and sometimes dangerous tasks for sailors and improving the efficiency of the service's assets.

With funding from the Office of Naval Research, the U.S. Naval Research Laboratory's (NRL) Naval Center for Space Technology (NCST) is building a robot that can perform routine tasks such as mopping decks, repainting, doing dishes, folding laundry, inspecting electrical panels and even looking for leaks. The assistance will allow sailors to focus more on the mission.

"All of the standard things that you need to do to have a ship operate, but those tasks are secondary to the ship's primary purpose, which is to fight naval battles," Glen Henshaw, head of robotics and machine learning in the lab's Naval Center for Space Technology, told GovCIO Media & Research. "The more that we can have robots do the routine maintenance tasks, the more we can have the sailors do the job that only sailors could do, and maybe have fewer sailors in harm's way when they do actually need to fight."

The lab initially wanted to build a small four-legged robot that would provide Marines with better intelligence, surveillance and reconnaissance capabilities, including inspecting a building without sending a human inside or navigating an outdoor terrain. But the Office of Naval Research became more interested in how the technology could be used aboard ships. Henshaw's team has been focusing on that aspect for the past three years.

Teaching a robot to interact with various objects around them effectively the way humans do with their hands every day without thinking about it is

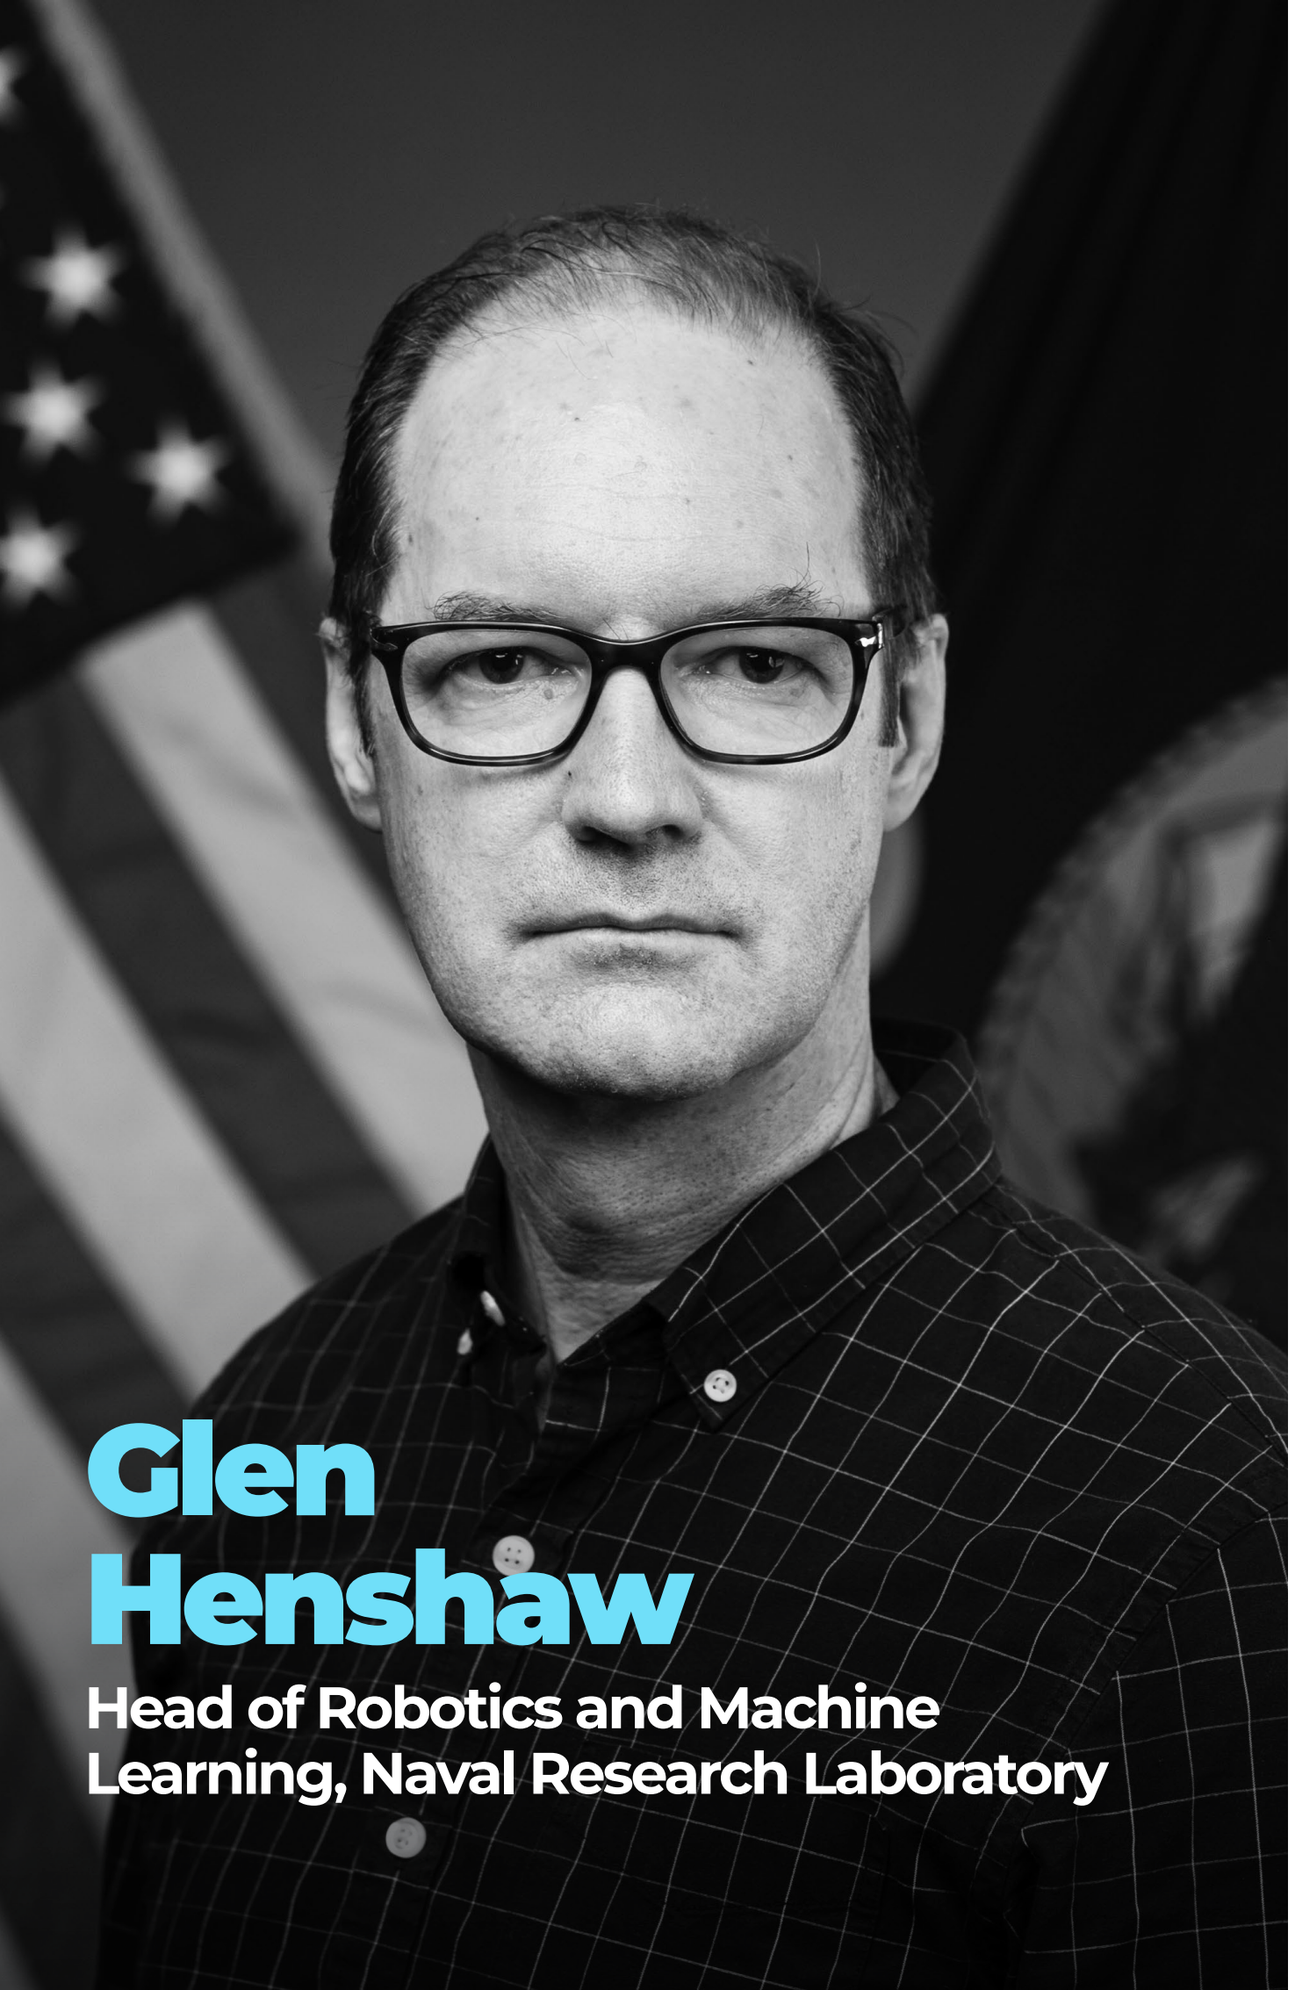

highly complex, and it is referred to in robotics as the “manipulation problem.”

“It’s very hard to build a robot that can do the dishes or fold the laundry. It’s hard enough just to get a robot that can open a door and walk through it, much less … handling objects that are really flexible or really fragile,” Henshaw said. “Those are the kinds of things that we would like robots to do onboard ships. And so we are applying types of machine learning to try to get robots to learn how to do those manipulation tasks.”

Solving the manipulation problem for the robot to be able to perform tasks such as grasping, lifting or placing items requires a combination of various processes, such as path planning, also known as motion planning and artificial intelligence techniques.

“[Robots] can’t really learn how to do motor tasks by us explaining them. They have to practice, so that’s a kind of AI called reinforcement learning. And we’re working very heavily with that,” Henshaw said.

Over the past 10 years, the open-source community has released high-quality software toolkits for AI, allowing laboratories interested in doing the type of work Henshaw’s team does to have a starting point with good software. Doing something innovative, however, requires writing more software on top of the software that’s publicly available.

“We try to embed software engineers in all of our groups because having a software engineer that both understands software, but also understands robotics and mathematics, and all of those sorts of specific issues to making software drive a device — those folks are really rare and very valuable,” Henshaw said.

To help potentially train AI developers, the Navy is looking to leverage generative AI systems and large language models without compromising data security.

“We have completed research on using large language models to convert Navy documents into executable agents,” David Aha, director of the Navy

Center for Applied Research in Artificial Intelligence, told GovCIO Media & Research. “Specifically, we used a fine-tuned version of GPT-3 to generate code from documentation samples, and evaluated these code samples for their quality and accuracy. We expect that this system could be used to tutor AI developers and, if given sufficient training samples, directly convert technical documents into executable agents.”

Another initiative is a collaboration with the Defense Advanced Research Projects Agency (DARPA) and Northrop Grumman called the Robotic Servicing of Geosynchronous Satellites (RSGS) program, which aims to build a robotic arm that repairs in-orbit satellites.

Currently planned for a 2024 launch, the program would revolutionize what can be accomplished in space. Satellites are fragile and expensive, and while some software issues can be easily resolved, most mechanical problems cannot be fixed, reducing their lifespan to only a couple of years.

“We’re finishing up the software, where we’re bolting everything together and very rigorously testing it to make sure that it’s going to work in space the way that it works on the ground,” Henshaw said.

For the lab, focusing on robotics and AI advancements overall has profound impacts.

“We’ll continue to work on AI and ML because it affects everything,” Steven Meier, director of the Naval Center for Space Technology, told GovCIO Media & Research.

“The more that we can have robots do the routine maintenance tasks, the more we can have the sailors do the job that only sailors could do, and maybe have fewer sailors in harm’s way when they do actually need to fight.”

— Glen Henshaw, Head of Robotics and Machine Learning, Naval Research Laboratory

GovCIO MEDIA & RESEARCH

DeepDives

DOD Responsible AI Tenets

It will require building a reliable governance structure and cultivating an educated and trained workforce to adopt AI responsibly at the DOD-wide levels.

1 RESPONSIBLE AI GOVERNANCE

Having clear guidelines and policies and ensuring disciplined governance structure all across DOD

2 WARFIGHTER TRUST

Building trust by providing training, education and establishing a test, evaluation, verification and validation framework

3 AI PRODUCT & ACQUISITION LIFECYCLE

Developing tools, policies, processes, systems and guidance for AI products throughout the acquisition lifecycle

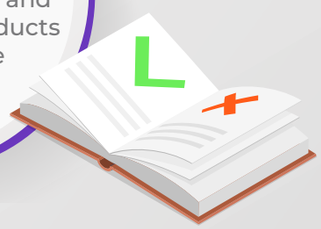

4 AI WORKFORCE

Building, training and retaining a RAI-ready workforce

5 RESPONSIBLE AI ECOSYSTEM

Building a responsible AI ecosystem on a global scale

6 REQUIREMENTS VALIDATION

Incorporating responsible AI into all applicable AI requirements

Source:
DOD Responsible AI strategy and Implementation Pathway

GovCIO MEDIA & RESEARCH

PARTNER INTERVIEW

THUNDERCAT TECHNOLOGY NetApp

AI Advancements Hinge on Data and Responsibility

Automation can have incredible impacts as long as the algorithms are built with key ethical practices in mind.

Kurt Steege
Chief Technology Officer, ThunderCat Technology

Caden Bradbury
AI/Analytics Solution Specialist, NetApp

What are the challenges you are seeing around artificial intelligence at DOD?

Steege Within DOD, the opportunities around AI, machine learning and deep learning are limitless. AI can play a part in cybersecurity, drone operations, tabletop exercises simulating deployment and enemy tactics, predictive maintenance, as well as many other basic needs like physical health, wireless coverage and even writing and maintaining contracts.

However, the same teams that are realizing these opportunities are faced with critical challenges. These include:

- **Data availability and quality:** The algorithms rely heavily on large amounts of quality data for training and making accurate predictions. DOD must be able to

“The agency is taking the correct approach of planning out a data platform first to allow AI systems to run optimally. Creating data lakes where multiple agencies can share data and insights will help to de-silo AI environments.”

— Caden Bradbury, Technical Solution Specialist, NetApp

efficiently acquire, clean, organize and share relevant data for AI applications.

- **Adversarial attacks and cybersecurity:** AI systems can be vulnerable to adversarial attacks where malicious actors intentionally manipulate input data to deceive or disrupt AI models.
- **Ethical and responsible use:** Data must come from a variety of sources to address concerns such as AI bias, transparency, accountability and the potential for autonomous systems to make life-or-death decisions.
- **Talent and expertise:** Recruiting and retaining AI experts, data scientists and engineers with relevant security clearances will be especially challenging.
- **Acquisition and procurement:** The rapid pace of development and complexity of AI systems both internally and in the cloud leads to challenges in acquisition, procurement and implementation.

What are some of the successes or use cases you’ve seen helping teams overcome some of these challenges?

Bradbury NetApp and NVIDIA implemented AI within the Navy by creating a joint solution for a project called Trident Warrior, which included the following components:

- NetApp Storage array
- NVIDIA Tesla V 100 GPUs
- 5G-connected cradle point
- Protopia AI’s machine-learning model
- NetApp’s Cloud Volumes ONTAP in Microsoft Azure
- Cameras to capture data

The goal of this project was to be able to detect bad actors on a Naval base. Cameras collected data of people’s faces and stored them in a small edge

device containing NetApp storage and NVIDIA compute. This data would then be used on Protopia AI's unique model to differentiate bad actors from those who are expected to be at the base. The key to this model is that the data is officiated, so the identity of those on camera is not understandable to people for security measures. This model and data could then be moved via a secure 5G network to Microsoft Azure, allowing for broader access and continued improvement to model and data quality.

Over the next year, what do you look forward to with regard to AI use within the government?

Bradbury DOD is looking at best practices from the private sector to design and scale AI operations. Building a fully optimized AI pipeline requires expertise in many areas, including data engineering, data science, enterprise applications, data center architecture, data visualization and domain knowledge. DOD will be looking to build a "dream team" of industry experts to help guide its AI journey. This wrangling will be a tall task, but we recognize DOD is focused on getting organized and ready to scale this year.

The process of sharing data across the department will also be top of mind. The agency is taking the correct approach of planning out a data platform first to allow AI systems to run optimally. Creating data lakes where multiple agencies can share data and insights will help to de-silo AI environments and significantly increase collaboration, leading to higher productivity.

AI applications perform best when data is as close to the AI operation as possible. Being able to move and replicate data across any platform (edge-near edge-core-cloud) is a game changer for the agency. Utilizing increasingly complex models across any environment is vital to AI success in DOD.

Photo Credit: Petty Officer 2nd Class Merissa Daley/DVIDS

THUNDERCAT TECHNOLOGY

NetApp

Adapt and Connect Seamlessly with ThunderCat Technology and NetApp AI

Learn More >

DOD CDAO Rethinks Adopting Department-Wide AI Acquisition Guidance

The AI office is considering whether to provide more guidance on AI acquisition.

BY ANASTASIA OBIS

As the Defense Department's Chief Digital and Artificial Intelligence Office is thinking through its acquisition best practices for artificial intelligence, officials said it might not develop a department-wide AI acquisition guidance after all.

"So, judging from my response, we are not far along in providing department AI acquisition guidance, but I don't know if we necessarily need to," the office's Deputy Chief Margaret Palmieri said at a RAND Corporation event.

"There's a testing evaluation piece around development that's not entirely acquisition. It could be, and I think we're trying to find that balance of … what degree is this a core set of tools that we ask people to use to what is, I hate to say it, a checklist, but to what extent is it a set of criteria that developers inside of government or industry must meet in order to field AI, whether that's on the development … or operational test and evaluation side?" she added.

Palmieri mentioned that the office already has its core acquisition vehicles, including Tradewind, a suite of services designed to accelerate the adoption of AI and machine learning, as well as data analytics solutions across the department. Last year, the platform launched the Tradewind Solutions Marketplace, a digital repository of post-competition designed to help the DOD solve its most pressing challenges around AI and machine-learning technologies.

"Tradewind is the platform really focused on kind of three things: increasing speed, enabling a variety of industry partners to play, and just agility in how well the contract and the needs can meet both the industry and the user, but not really on the path for policy quite yet," Palmieri said.

"What we've really been trying to wrestle with is not to over-centralize because the department is so diverse and distributed and so large that we want innovation to happen at the edge," she added. (ctd.)

U.S. Air Force Tech. Sgt. Alyssa Wier, weapons director assigned to the 176th Air Defense Squadron, Alaska Air National Guard tests the new Battle Management Training NEXT system Aug. 26, 2021.

Photo credit: Maj. Kimberly Burke / DVIDS

The office has also been exploring how it can apply the large language models to defense use cases. It has been experimenting with different generative AI models through a series of experiments called the Global Information Dominance Experiments (GIDE) to test solutions around AI processes.

“Really, just to test out, you know, how do they work? Can we train them on DOD data, or tune them on DOD data? How do our users interact with them? And then what metrics do we want to come up with based off of what we were seeing to facilitate evaluation of these tools? Because there are really great evaluation metrics for generative AI,” Palmieri said.

Palmieri emphasized that DOD needs to pay more attention to the possible negative consequences of the technology, particularly referring to “hallucination” when the technology generates false information.

“There are going to be use cases that it’s really really good for, and there are going to be use cases it’s not good for. What we found is there’s not enough attention being paid to the potential downsides of generative AI — specifically, hallucination,” Palmieri said

“This is a huge problem for DOD and it really matters for us and how we apply that, and so we’re looking to work more closely with industry on these types of downsides and not just hand-wave them away,” she added.

“There are going to be use cases that it’s really really good for, and there are going to be use cases it’s not good for. What we found is there’s not enough attention being paid to the potential downsides of generative AI — specifically, hallucination.”

— Margaret Palmieri, Deputy Chief Digital and Artificial Intelligence Officer, DOD