OCTOBER 2020

U.S. Military Investments in Autonomy and AI

A Budgetary Assessment

CSET Policy Brief

AUTHORS

Margarita Konaev
Husanjot Chahal
Ryan Fedasiuk
Tina Huang
Ilya Rahkovsky

CENTER for SECURITY and EMERGING TECHNOLOGY

Executive Summary

National security leaders identify artificial intelligence as a priority technology for defending the United States. Today's investments in research and development across the federal government, academia, and industry will help secure America's global leadership in AI. Investments in defense research in particular will provide the U.S. military with the AI-enabled capabilities needed to deter adversaries from aggression, fight and win the wars of the future, and cooperate effectively with allies. But where exactly is this investment going? And what benefits and risks might result from developing and fielding autonomous and AI-enabled weapons and systems?

Policymakers need information about the Department of Defense's investments in AI for proper oversight and to ensure these research efforts support broader strategic goals. As the U.S. defense community implements its vision for AI, CSET offers a two-part analysis assessing the scope and implications of U.S. military investments in autonomy and AI.

Drawing on publicly available budgetary data on DOD's science and technology (S&T) program and an extensive review of strategic and operational literature and scientific research, these studies focus on three interconnected elements that form our analytical framework:

- The *technology* element addresses DOD research and development efforts in autonomy and AI;
- The *military capabilities* element speaks to the speed, precision, coordination, reach, persistence, lethality, and endurance enabled by advances in autonomy and AI;
- The *strategic effects* element analyzes how these technological developments and capability enhancements may affect key strategic issues—specifically, deterrence, military effectiveness, and interoperability with allies.

This report centers on the technology element, while the accompanying report, "U.S. Military Investments in Autonomy and AI: A Strategic Assessment," covers the military capabilities and strategic effects portions. The following is a summary of our findings regarding current DOD research investment priorities, trends, and gaps with corresponding recommendations.[1]

Current DOD Research Investments: Trends and Gaps

The U.S. military has a wide range of research programs using autonomy and AI in unmanned vehicles and systems, information processing, decision support, targeting functions, and other areas. Yet there are gaps in research on AI not related to autonomy and in investments in basic AI research.

Our results show that estimates of research investments vary depending on definitions and measures. Across the different measurements, however, the data suggests that:

- AI research unrelated to autonomy—and especially autonomy in unmanned systems—receives a relatively small share of the S&T funds directed toward autonomy and AI research.
    - More than two-thirds of the funds allocated to AI-related science and technology research are also related to autonomy.
    - In contrast, less than a third of the funds allocated to autonomy-related research were related to AI.
- Investments in basic AI research are also likely smaller than initially estimated.
    - Three programs—DARPA's "Defense Research Sciences Program," the Air Force's "University Research Initiatives," and the Navy's "Defense Research Sciences" program—account for about 80 percent of the funds allocated to basic AI research. The estimated program-level costs of these initiatives, however, are highly inflated because they include the costs of research projects unrelated to AI.

The ambiguity about the nature and scope of U.S. military investments in autonomy and AI research makes it difficult to ensure oversight. Moreover, the current U.S. military research on AI may not be sufficiently innovative to fuel the scientific breakthroughs needed to ensure long-term advantage. As such, we offer the following policy recommendations:

- DOD should provide greater clarity about overall funding levels for autonomy and AI, overlap between funding allocated to autonomy research and AI research, and funding for AI-related basic research.
- DOD should leverage its relationships with university-affiliated research centers and national labs to map the landscape of non-

autonomy related AI and potential military applications, and identify opportunities for additional investment.

Effective human-machine collaboration is key to harnessing the full promise of AI. But gaps in our understanding of trust in human-machine teams can impede progress.

The U.S. military sees many benefits to pairing humans with intelligent technologies and our analysis finds that human-machine collaboration is a crosscutting theme across the different autonomy and AI research programs. The following issues therefore merit attention:

- Trust is essential to human-machine collaboration. Yet in our assessment, few autonomy and AI related research initiatives reference both trust and human-machine collaboration.
  - Only 18 out of the 789 research components related to autonomy and 11 out of the 287 research components related to AI mention the word "trust."
- Gaps in research on the role of trust in human-machine teams can negate the advantages in speed, coordination, and endurance promised by autonomy and AI. This, in turn, could impede U.S. ability to use AI-enabled systems to deter adversaries from aggression, operate effectively on future battlefields, and ensure interoperability with allies.

To safely and effectively employ machines as trusted partners to human operators, the following steps may be necessary:

- DOD should increase investment in multidisciplinary research on the drivers of trust in human-machine teams, specifically under operational conditions.
- DOD should assess the advantages of making trust a consistent theme across autonomy and AI research programs pertaining to human-machine collaboration.

Table of Contents

Introduction

The Department of Defense has an ambitious vision for AI. And while most federal agencies have seen their research and development funding decline in 2020, DOD's R&D budget has increased.[2] The FY2021 U.S. defense budget request allocates $1.7 billion to autonomy to enhance "speed of maneuver and lethality in contested environments" and the development of "human/machine teaming," as well as $800 million to AI, building on previous funds directed to the Joint Artificial Intelligence Center (JAIC) and Project Maven.[3] These investments are both timely and welcome. Yet as the Department of Defense Artificial Intelligence Strategy warns, other nations—particularly China and Russia—are also investing in military applications of AI, threatening to erode U.S. "technological and operational advantages and destabilize the free and open international order."[4]

If leveraged correctly, investments in autonomy and AI research will help the United States field a more lethal and resilient military, make faster and better decisions in crises and wartime, and coordinate forces and assets in hostile and denied environments. The overarching objective of our research was to better understand the Department of Defense science and technology investments in autonomy and AI, and how they may affect the military capabilities and strategic interests of the United States. Toward that end, we developed the following analytical framework, which focuses on three interrelated elements:

Technology: Mapping the landscape, type, and monetary value of U.S. military research and development programs related to autonomy and artificial intelligence, particularly in areas like unmanned systems, information processing, decision support, and the targeting process.

Capabilities: Assessing the critical capabilities enabled by advances in autonomy and AI, with specific attention to speed, precision, coordination, reach, persistence, lethality, and endurance.

Strategic Effects: Analyzing how these technological developments and capability enhancements may affect U.S. strategic objectives, specifically, deterring adversaries, conducting effective military operations in line with the Law of War, and ensuring interoperability with allies.

Figure 1. Analytical framework

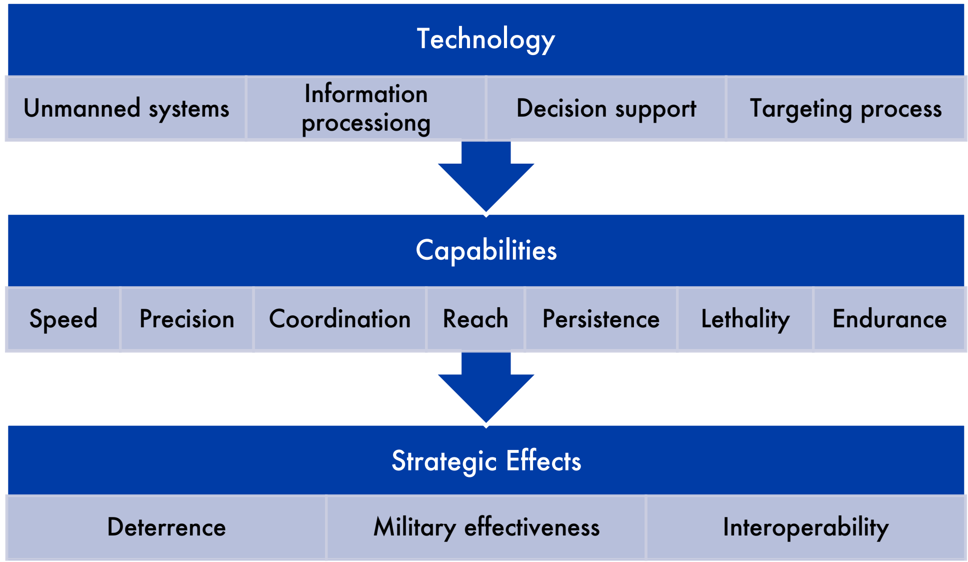

This report is the first of two CSET studies on the scope and implications of U.S. military investments in research and development related to autonomy and AI. Here, we focus on the technology component of our tripartite analytical framework, presenting the findings from a budgetary assessment of DOD's FY2020 S&T program investments, based on publicly available budgetary data. The second report, "U.S. Military Investments in Autonomy and AI: A Strategic Assessment," elaborates on the military capabilities and strategic effects elements.

U.S. Military S&T Investments in Autonomy and AI: A Budgetary Assessment

America's global competitiveness in AI is fueled by scientific and technological breakthroughs in federal, academic, and commercial research and development.[5] Defense research, however, explicitly strives to ensure the U.S. military has the speed, precision, coordination, reach, persistence, lethality, and endurance necessary to deter adversaries, fight and win the wars of the future, and cooperate with allies.

Our analysis of U.S. military investments in autonomous and AI-enabled weapons and systems draws on the FY2020 research, development, testing, and evaluation (RDT&E) budget justification books of the Army, Air Force, Navy, Marines, and DARPA. We focus specifically on basic, applied, and advanced research—known jointly as the Science and Technology program, which supports the development of new technologies imperative to U.S. military superiority.

After reviewing the methodology, we estimate the financial resources dedicated to S&T programs related to autonomy and AI, highlight specific focus areas for these technologies and applications, and evaluate whether these investments align with broader U.S. strategic goals.

Methodology

There is no consensus on the definition of autonomy or AI in the scientific literature or in the national security community. Appendix I provides a range of definitions from pertinent government, technology, and national security sources. In this report, we adopt an approach (described in more detail in this section) that is broadly inclusive of what could be considered AI.

Deciphering the relationship between autonomy, AI, and related terms such as robotics and automated systems can be difficult; these terms are often used interchangeably, and existing definitions can conflict with one another.[6] The Defense Science Board Summer Study on Autonomy offers a helpful differentiator between automated systems, which are "governed by prescriptive rules that permit no deviation," and autonomous systems, which "must have the capability to independently compose and select among different courses of action to accomplish goals based on its knowledge and understanding of the world, itself, and the situation."[7] However, because no system is truly autonomous, autonomy is generally understood as a capability or in reference to autonomous functions.

The 2018 Department of Defense Artificial Intelligence Strategy defines AI as "the ability of machines to perform tasks that normally require human intelligence—for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action—whether digitally or as the smart software behind autonomous physical systems."[8] DOD has developed and deployed AI technologies for more than 40 years. While early systems were largely automated, recent advances in AI enable greater autonomous functionalities in intelligent systems—both those operating virtually in software (e.g., autonomy at rest) and those with a presence in the physical world, such as robotics and autonomous vehicles (e.g., autonomy in motion).

Our analysis of U.S. military investments in basic, applied, and advanced research related to autonomy and AI uses data collected from the aforementioned FY2020 RDT&E budget justification books. These documents provide Congress with detailed information about the estimated costs, duration, rationale, and intended impact of ongoing and planned programs.[9] Basic research focuses on gaining an understanding of fundamental aspects of a given phenomenon, applied research focuses on application-specific knowledge, and advanced research looks toward field-testing and integration of hardware.[10]

The information captured in each document is organized in three tiers: a) program elements (PE), b) projects, and c) what we call "components." PE is the largest unit comprising projects, and each project further consists of components. At each tier, the data we collected includes a description, associated names and/or numbers, and budget allocations ranging from FY2018 to FY2024.

Naturally, not all the programs detailed in the justification books are related to autonomy and/or AI. To identify those focusing on autonomy and/or AI, we conducted a keyword search scanning the descriptions of PEs, projects, and components. Although not an exhaustive list, the keywords presented in Table 1 encompass a range of research methods, techniques, applications, systems, and functions related to autonomy and AI.

Table 1. Autonomy and AI keywords

Autonomy Keywords:
Algorithm, automate, automatic, automation, autonomous, autonomy, cybernetics, fully autonomous, human operator, human-agent interaction, human-agent teaming, human-control, human-in-the-loop, human(-|/)intelligent agent, human-intelligent team, human-machine, human-machine collaboration, human-machine interface, human-machine teaming, human-on-the-loop, human-out-of-the-loop, human-robot, human-supervised, human-system, human/unmanned, intelligent agent, intelligent control, intelligent system, machine-control, manned(-|/)unmanned teaming, optionally manned, predictive, robot, semi-autonomous, soldier-intelligent agent, supervised-autonomous, swarm, uninhabited, unmanned.

AI Keywords:
Active learning, adaptive learning, anomaly detection, artificial intelligence, associative learning, autonomous navigation, autonomous system*, autonomous vehicle*, average link clustering, back propagation, backpropagation, binary classification, bioNLP, boltzmann machine, character recognition, classification algorithm, classification label*, clustering method*, complete link clustering, computer aided diagnosis, computer vision, deep learning, ensemble learning, evolutionary algorithm, fac* expression recognition, fac* identification, fac* recognition, feature extraction, feature learning, feature matching, feature selection, feature vector, feedforward network, feedforward neural network, fuzzy clustering, generative adversarial network, gradient algorithm, graph matching, graphical model, handwriting recognition, hierarchical clustering, hierarchical model, human robot, image annotation, image classification, image matching, image processing, image registration, image representation, image retrieval, incremental clustering, information extraction, information fusion, information retrieval, k nearest neighbor, knowledge based system*, knowledge discovery, knowledge representation, language identification, machine learning, machine perception, machine translation, multi class classification, multi label classification, multi task learning, natural language

generation, natural language processing, natural language understanding, neural network, object recognition, one shot learning, pattern matching, pattern recognition, random forest, recommend* system*, recurrent network, reinforcement learning, scene* classification, scene* understanding, semi supervised learning, sentiment classification, single link clustering, spatial learning, speech processing, speech recognition, speech synthesis, statistical learning, strong artificial intelligence, supervised learning, support vector machine, text mining, text processing, transfer learning, translation system, unsupervised learning, video classification, video processing, weak artificial intelligence, zero shot learning.

Source: List of keywords compiled by CSET researchers.

To the degree possible, we distinguish between autonomy-focused and AI-focused lines of research to better understand the different capabilities, functions, and tasks DOD expects each of these technologies to enable and perform. The emphasis on unmanned systems and human-machine interactions in the autonomy keywords list corresponds to U.S. defense budget language for FY2020 and FY2021.[11] The AI keywords list covers different fields of research related to AI, ranging from theoretical models to emerging applications and mature technologies—all gleaned from a variety of research publications.[12]

Using the justification books offers a number of benefits. First, they provide the most reliable and consistent open-source resource on DOD RDT&E investments in autonomy and AI. Second, they reflect DOD priorities and assessments of which technologies and capabilities the United States needs to advance its interests. Third, these documents provide insights into the domains, operational environments, warfighting functions, and tasks for which the department envisions leveraging autonomy and AI.

The justification books also have a number of limitations. First, the PE, project, and component descriptions lack the detail needed to fully understand the technical dimensions of each program. Second, given the heightened interest in advancing AI capabilities throughout the defense community, individual program managers may have included AI or related keywords in descriptions of programs in which AI is not a substantial component. Third, each service has their own style, leading to variation in the level of program detail offered. Basic, applied, and advanced research programs are also discussed in different ways.[13] Together, these factors could lead us to underestimate or overestimate the scope of investments in autonomy and AI-related programs.

To overcome these challenges, we developed a comprehensive list of keywords for autonomy and AI (see Table 1) and conducted a series of robustness checks to prevent systematic bias in our estimates. We provide a range of calculations including component level, PE level costs, and maximum estimates stratified by research category and service. While moderately confident in the cost calculations we present, the limitations of these estimates are discussed in the following section. We encourage thinking of these investments as an indicator of priorities rather than purely in financial terms.

Budgetary Analysis

The FY2020 RDT&E budget justification books of the Army, Air Force, Navy, Marines, and DARPA include a total of 143 PEs, 674 projects, and 2,312 components across the basic, advanced, and applied research categories. Using the autonomy and AI keywords in Table 1 to query the descriptions of PEs, projects, and components offered by these documents, we classified 44 PEs (31 percent), 201 projects (30 percent), and 789 components (34 percent) as related to autonomy, and 13 PEs (9 percent), 62 projects (9 percent), and 287 components (12 percent) as related to AI.

Calculating the exact amount of money the U.S. military allocates to research related to autonomy and AI based on this data is an intricate process. For instance, one could focus on budget estimates at the PE level because they include project- and component-level estimates. However, this approach has at least two problems. First, PEs classified as related to autonomy and/or AI also include lines of research at the project and component level ***not*** related to autonomy or AI. Thus, by focusing exclusively on PE level costs, we would certainly overestimate the amount of money directed to autonomy and AI research in PEs classified as autonomy and/or AI related. Second, in our query, we have identified projects and components related to autonomy and/or AI clustered under PEs that were ***not*** classified as such. Therefore, focusing only on the autonomy and/or AI related PEs risks undercounting these other efforts.[14]

Alternatively, one could focus on autonomy or AI related components. The budget figures here are much closer to actual investments in autonomy and AI related research than PE-level estimates.[15] That said, unlike the PE- and project-level costs—available for FY2018–FY2024—the component level costs are only publicly available for FY2018–FY2020. We therefore cannot use this measure to estimate the budget for new lines of research. Moreover, not all projects disaggregate cost at the component level. Thus, though arguably more accurate than PE-level estimates, the total sum of costs at the component level likely underestimates investments made in autonomy and AI research.

Considering these challenges, we offer three estimates of U.S. military S&T investments in autonomy and AI:

- **Component-level estimates** calculated as the sum of the costs at the component level.
- **PE-level estimates** calculated as the sum of the costs at the PE level.
- **Maximum estimates** calculated as the sum of the three following measures:
    - the cost of PEs classified as related to autonomy or AI
    - the cost of projects classified as related to autonomy or AI derived from PEs ***not*** related to autonomy or AI, and
    - the cost of components classified as related to autonomy or AI derived from projects ***not*** related to autonomy or AI.

We are moderately confident in our estimates. Nonetheless, given the aforementioned challenges and the different approaches to defining and measuring autonomy and AI, we caution the reader about possible discrepancies between our calculations and other sources and across the figures we provide.

Estimating the value of U.S. military S&T investments in autonomy and AI

Tables 2 and 3 present the component, PE, and maximum estimates of U.S. military S&T investments in autonomy and AI, stratified by service and research category, respectively.

Table 2. Component, PE, and maximum estimates of U.S. military S&T budget related to autonomy and AI, by service (USD in millions, FY2018–FY2020)

Autonomy			
	Component	PE	Maximum
Army	2,873	3,982	5,735
Navy	4,866	1,885	5,345
Air Force	2,861	1,026	3,918
DARPA	4,938	7,437	8,446
AI			
	Component	PE	Maximum
Army	1,194	896	2,227
Navy	2,465	1,417	3,655
Air Force	642	481	1,482
DARPA	1,586	3,110	4,200

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Table 3. Component, PE, and maximum estimates of U.S. military S&T budget related to autonomy and AI, by research category (USD in millions, FY2018–FY2020)

Autonomy			
	Component	PE	Maximum
Basic	3,080	2,676	4,075
Applied	7,320	7,066	10,931
Advanced	5,139	4,589	8,438
AI			
	Component	PE	Maximum
Basic	1,727	3,157	3,880
Applied	2,878	2,547	5,587
Advanced	1,281	201	2,098

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

The evidence presented in Tables 2 and 3 highlights an important point:

- U.S. military S&T investments in autonomy and AI vary significantly based on the method of calculation. The comparison between component-level costs, PE-level costs, and maximum estimates illustrates the shortcomings of a singular estimate of U.S. military spending on research and development related to autonomy and AI.[16]

Policymakers and national security experts are interested in how much the United States military invests in emerging technologies. Many worry that high Chinese government spending on AI R&D may outflank the United States in the global competition for AI leadership.[17] Congress and other regulatory agencies must also know the financial scope of these investments to ensure accountability and oversight.

However, as demonstrated here, no easy method exists for accurately estimating U.S. military research and development investments in autonomy and AI. The amounts vary significantly depending on the calculation technique and different definitions of complex, overlapping terms such as autonomy and AI. Still, two clear themes emerge: 1) AI research unrelated to autonomy—and especially autonomy in unmanned systems—receives a relatively small share of the S&T funds directed toward autonomy and AI research, and 2) investments in basic AI research are likely lower than initially estimated. We discuss each in turn below.

Most AI-related research efforts are also related to autonomy

AI is often considered the intellectual foundation of autonomy, and AI applications can enable autonomous functionalities in physical systems or robots. Autonomy, however, is not the same as AI.[18] Rather, it is a broad discipline defined differently depending on context or purpose.

Our analysis of U.S. military S&T investments in autonomy and AI shows that the majority of AI related research efforts across the three tiers of PE, project, and component were also related to autonomy: 9 of 13 AI PEs (69 percent), 36 of 62 AI projects (58 percent), and 198 of 287 AI components (69 percent) were also classified under autonomy. In contrast, the majority of autonomy research efforts were not classified as related to AI, with 35 of 44 PEs (80 percent), 165 of 201 projects (82 percent), and 591 of 789 components (75 percent) identified as related to autonomy but not AI. In other words:

- While most AI related research efforts are also related to autonomy, the majority of autonomy research efforts are not classified as AI.

That autonomy-related research constitutes such a large part of the U.S. military investment in AI research is also reflected in the budget allocations. Based on component-level estimates, Figure 2 shows the share of the autonomy research budget that is also classified as AI and the share of the research funds allocated to AI that were also classified as related to autonomy.

Figure 2. U.S. military S&T budget funds allocated to autonomy and AI, by service and research category (component-level estimates, FY2018–FY2020)

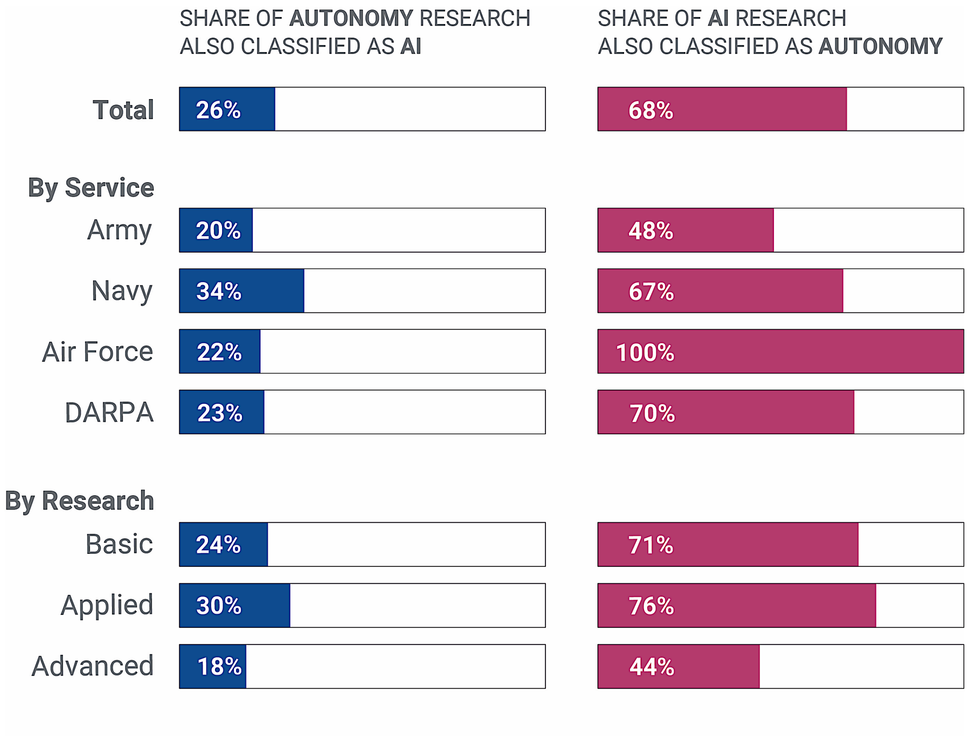

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Looking at component-level estimates, 48 percent of the Army's research funds, 67 percent of the Navy's, 70 percent of DARPA's, and 100 percent of the Air Force's research funds allocated to AI were also classified as related to autonomy. Estimating this distribution by research category (also at the component level), 71 percent of basic research funds, 76 percent of applied research funds, and 44 percent of advanced research funds allocated to AI were also classified as related to autonomy.

Figure 3 illustrates the overlap between research efforts related to autonomy and those identified as related to AI, looking at both the number of AI components and the funds dedicated to researching these technologies.

Figure 3. Most of the U.S. military S&T research related to AI also related to autonomy (component-level estimates, FY2018–FY2020)

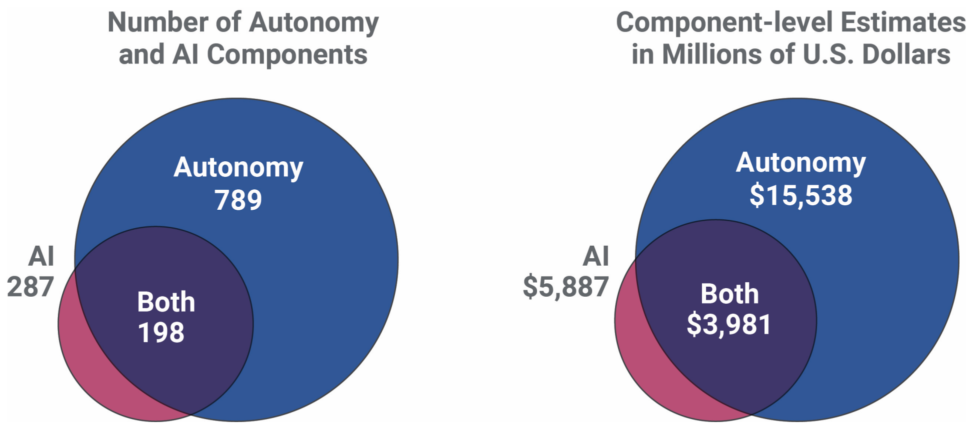

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

As evident from the above analysis, much of the U.S. military AI research portfolio is intertwined with its autonomy research, and autonomy research is skewed toward unmanned systems. A large portion of the U.S. military's AI research is therefore likely geared toward AI applications that enable greater autonomous functionalities in unmanned systems, such as robotics and different semi-autonomous and autonomous vehicles. It follows that:

- The amount of money and share of funds dedicated to AI research unrelated to autonomy and unmanned systems is likely smaller than it appears.

The potential gap in non-embodied AI research is concerning, given the proliferation of malicious cyber operations and the vulnerability of critical infrastructure. Additional focus on AI-enabled cyber defense research, for instance, could be timely.

Gaps in investment in basic AI research

At a first glance, U.S. military investment in basic AI research appears relatively high. Basic AI research was allocated between $7.4 and $9.1 billion from FY2018 to FY2024. Going by PE-level estimates, basic AI research received more funds than applied research or advanced research. (Annual estimates are provided in the Appendix.)

There are, however, two reasons to believe that actual investment in basic AI research is much smaller than the figures above indicate. First, of the funds broadly allocated to basic AI research at the PE level, only a fraction is directed toward research that explicitly focuses on AI. There are three PEs classified as related to AI in basic research: DARPA's "Defense Research Sciences Program," the Air Force's "University Research Initiatives," and the Navy's "Defense Research Sciences" program. These PEs represent some of the most expensive lines of effort related to AI. Yet a closer look at the component-level allocations under these PEs shows much lower investments in AI specifically:

- The estimated cost of DARPA's "Defense Research Sciences" PE is $2.9 billion for FY2018–FY2024. Two of the six projects and 19 of the 50 components in this PE relate to AI. At the component level, AI-related research efforts are estimated at $483 million for FY2018–FY2020.[19]

- The estimated cost of the Air Force's University Research Initiatives PE is $1.15 billion for FY2018–FY2024. There is only one project under this PE and it is coded as related to AI, but none of the six components had AI-related keywords in their descriptions. The program supports "defense-related basic research in a wide range of scientific and engineering disciplines relevant to maintaining U.S. military technology superiority," with research topics including "transformational and high priority technologies such as nanotechnology, sensor networks, artificial intelligence" and many others.[20] While it is difficult to ascertain precisely what portion of the funds under this PE is directed to AI research, the description suggests AI is only one of many priorities.

- The estimated cost of the Navy's Defense Research Sciences PE is $3.4 billion for FY2018–FY2024. Neither of its two projects, and only eight out of the 15 components were classified as related to AI. At the component level, AI-related research efforts are estimated at $846 million for FY2018–FY2020.[21]

Notably, component-level estimates only cover FY2018–FY2020, while project and PE level numbers relate to FY2018–FY2024. That said, the discrepancy between the PE-level estimates and component-level estimates illustrates that only a fraction of the research funds allocated under DARPA's "Defense Research Sciences Program," the Air Force's "University Research Initiatives," and the Navy's "Defense Research Sciences" programs goes directly toward basic AI research.

Second, as previously noted, funding for AI is often allocated in projects and components under PEs not classified as AI (e.g., non-AI PEs). Thus, to get a holistic picture of the funds directed toward AI research, it is important to account for these estimates as well. Figure 4 shows how AI-related research funds are distributed across the different research categories—basic, applied, and advanced research—paying attention to both the funds allocated under the AI PEs, as well as the money found in AI-related projects and components nested under non-AI PEs.

Figure 4. Distribution of AI-related research funds by research category: PE-level estimates and funds allocated under non-AI PEs (USD in millions, FY2018–FY2024)

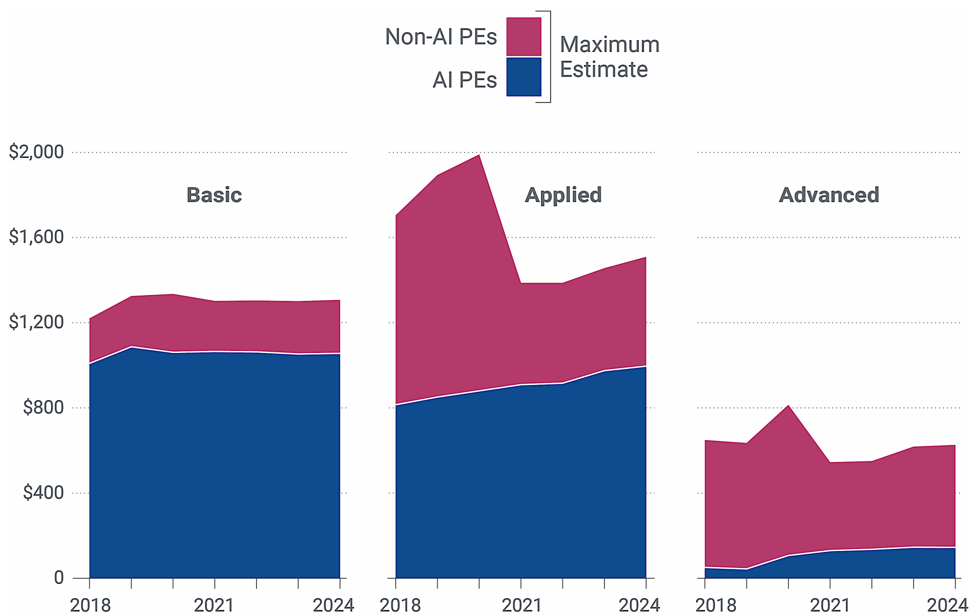

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

As Figure 4 demonstrates, there are notable differences in how funding is distributed across the different research categories. For basic research, most of the money is found under AI PEs; for applied research, the funding allocated under AI PEs and non-AI PEs is approximately equal; and for advanced research, most of the money is classified under non-AI PEs. This discrepancy is important when assessing the true value of funding allocated to basic AI research.

Now, as Figure 4 shows, the share of funds allotted to basic AI research in projects and components nested under non-AI PEs is much smaller than what is allocated under AI PEs. In fact, approximately 80 percent of the funds allocated to basic AI research are designated under the three aforementioned AI PEs: DARPA's "Defense Research Sciences Program," the Air Force's "University Research Initiatives," and the Navy's "Defense Research Sciences" program. Yet as previously discussed, the PE-level estimates of these programs are highly inflated because they include the cost of projects and components unrelated to AI.

Thus, while the money found in projects and components nested under non-AI PEs amounts to less than 20 percent of the funds dedicated to basic AI research, the real value of the investment in basic AI research catalogued under the three aforementioned AI PEs—namely, the remaining 80 percent—is much smaller than the PE-level estimates suggest. It follows that the actual investment in basic AI research is likely much smaller than initial estimates indicate.

This is concerning because investments in long-term fundamental research in AI and ML are necessary for breakthroughs that could deliver a lasting U.S. advantage against strategic competitors and adversaries. That said, CSET research suggests Chinese public AI R&D spending focuses heavily on applied research and experimental development, not basic research.[22] An opportunity may yet exist to gain an advantage through investments in areas neglected by U.S. competitors.

Areas of focus

We offer a snapshot of U.S. military basic, applied, and advanced research and development initiatives related to autonomy and AI in four broad categories: unmanned systems, information processing, decision support, and targeting functions. Analysts classify defense research on emerging technologies in a myriad of ways. Some discuss developments in these areas through the lens of the mission—for instance, advances in AI for cyberspace operations or information operations.[23] Others classify AI applications by the nature of the environment where they may be used, drawing a line between enterprise AI, mission-support AI, and operational AI.[24] Others still focus on functions—for example, autonomous functions for mobility, targeting, intelligence, and more.[25]

There is no single framework that can perfectly catalogue these efforts. Autonomy and AI are general purpose technologies with potential applications across many different fields, and based on their descriptions, certain S&T initiatives are sufficiently broad to be classified under all four categories we review. Moreover, our data encompasses research and development efforts at different levels of maturity. The following discussion therefore illustrates the range of functions, tasks, mission areas, and operational contexts for which the U.S. military envisions using autonomy and AI. It is not comprehensive assessment of all the different research and development lines of effort advanced by the U.S. military.

Unmanned systems: Numerous lines of effort across the services are dedicated to increasing automation, autonomy, and AI in unmanned aerial vehicles and systems, unmanned ground vehicles and systems, unmanned undersea vehicles, and unmanned surface vehicles. AI technologies in this space are used to perceive and map the environment, fuse sensor data, identify obstacles, plan navigation, and communicate with other vehicles. Unmanned weapon systems use AI, as well as semi-autonomous and autonomous capabilities, for functions such as detecting, identifying, selecting, tracking, or engaging targets. Unmanned and unarmed military systems with autonomous and AI-enabled functions can be used for intelligence, surveillance, and reconnaissance (ISR), logistics, resupply, and maintenance.

- The Army's "Next Generation Combat Vehicle technology" is one the main programs using autonomy and AI in unmanned systems for a broad range of functions. Its "Artificial Intelligence and Machine Learning Technology" project, for instance, develops AI and ML software and algorithms to "team with soldiers in support fully autonomous maneuver of NGCV and other autonomous systems, both physical and non-embodied." The project develops capabilities for NGCV that will "increase autonomy, unburdening the soldier operator, with a high degree of survivability and lethality in a highly contested environment."[26]

- The NGCV Advanced Technology program matures and demonstrates some of the efforts at the applied level, including using AI/ML techniques "to develop/integrate intelligent formation control to be used on maintained roads and in complex terrain without the need for GPS."[27]

Information processing: Advances in artificial intelligence—such as high-fidelity sensing, machine learning, computer vision, and natural language processing—allow systems to collect, collate, and analyze complex data at unprecedented speed and volumes. AI applications for information processing can free up personnel for higher-order tasks. AI-enabled speed and accuracy in information processing also improves situational awareness and helps leaders at all levels of command make better decisions. Across the services, there are many research programs on AI for information processing in support of different military functions and missions.

- The Air Force applied research project on "Sensor Fusion Technology," for instance, explores "concepts and algorithms for efficient parallel processing, distributed processing, and high-performance computing in sensor data processing and synthetic data generation" in support of ISR, situational awareness, battlefield visualization, target recognition, and battle damage assessment capabilities against a wide variety of targets.[28]

- DARPA's "Automated Knowledge Acquisition" program develops technologies that "automatically learn the semantics of a new data source, characterize source content, align source schema to the target, transform and load values, and reconcile inconsistencies by learning from previously integrated sources." AKA technologies aim for human-free data integration to "automatically create and maintain, in real-time, broad knowledge of local and regional military, political, economic, social, and cultural information for warfighters in theater."[29]

Decision support: Military decision-making processes and procedures are often time- and resource-intensive. Advances in big data analytics, speech recognition systems, natural language processing, neural networks, reinforcement learning, and other techniques will help commanders process and assess more options for action in complex situations.[30] Decision support technologies enable a shift from situational awareness to situational understanding; from data to information to knowledge; and ultimately, from real-time analysis to prediction and automatic proposal of action plans. Some of the U.S. military research efforts applying AI to decision support include:

- Navy's "Applied Research Innovative Naval Prototype" will create AI applications for "predictive mission-focused analytics that autonomously gather, analyze, compile, interpret, and visualize a fused tactical & national all source data picture to improve decision making speeds and enable a distributed Artificial Intelligence capability that can function in a harsh and adversarial environment and is able to determine an optimal response and accelerate reactions to real time."[31]

- Army basic research "Data-to-Knowledge to Support Decision Making" includes "research in support of rapidly enhancing long-duration, complex, dynamic decision-making capabilities of individual Warfighters and units through the integration of cognitive augmentation and course of action recommender technologies."[32]

- Army's applied "Expeditionary Data to Decisions Technology" research examines "artificial intelligence techniques that provide the most relevant and available data to support time-sensitive and critical decisions, and present information in context and in alignment with complex cognitive needs."[33]

Targeting functions: At least 154 weapon systems around the world already use autonomous functions to support aspects of the targeting process—from identification, tracking, prioritization, and selection of targets to, in some cases, target engagement.[34] Current U.S. military S&T initiatives in this area include:

- Research on Automatic Target Recognition (ATR):
    - Army basic "ATR Research" seeks to provide the capability to predict, explain, and characterize the signature content of different targets and background to reduce the workload on the analyst and enhance the robustness and effectiveness of land warfare systems.[35]
    - DARPA's "ATR Technology" research focuses on three core areas: "(1) development of on-line adaptive algorithms that enable performance-driven sensing and ATR utility; (2) algorithm training technology that enables rapid incorporation of new targets; and (3) technologies that dramatically reduce required data rates, processing times, and the overall hardware and software demands of ATR systems."[36]
- Initiatives incorporating targeting research within broader research efforts on unmanned systems, such as:
    - Navy applied research on undersea weaponry focuses on developing technologies for unmanned undersea vehicles including enablers for pre-engagement positioning and fire-control solutions for effective weapon-to-target engagements.[37]
    - Army's applied "Next Generation Combat Vehicle" program has a "Next Generation Intelligent Fire Control Technology" project that "investigates image sets for computer vision

algorithms, target acquisition validation schemes and experimentation of large caliber armament systems."[38]

- Army's advanced "Future Vertical Lift Advanced (FVL) Technology" program includes research that applies AI to full spectrum targeting, looking to "mature and demonstrate key targeting sensor system and automation (i.e., Artificial Intelligence/Machine Learning) technologies essential to enable the FVL and Future Unmanned Aircraft System modernization priorities."[39]

A crosscutting focus on Human-Machine Collaboration and Teaming: Public discourse on military AI often invokes scenarios where machines replace humans on future battlefields and wartime decisions are made without human control. Our assessment of U.S. military research on emerging technologies points to a different vision of future warfare: artificial intelligence will augment human intelligence, as humans and intelligent systems interact, collaborate, and integrate to optimize performance both for humans and human-machine teams in different environments and missions.

Human-machine collaboration is a broad term encompassing the many ways in which humans and intelligent technologies—robots, intelligent agents, unmanned systems, AI-enabled autonomous systems, and other non-human intelligent agents—interact. The U.S. military sees many benefits to pairing humans with machines, including reducing risk to U.S. service personnel, lightening the cognitive and physical load to improve performance and endurance, and increasing accuracy and speed in decision-making and operations. Today's AI-enabled technologies are more tools than teammates, and are largely unprepared for safe operational deployment, especially in combat conditions. But in the future, advanced human-machine teaming may feature intelligent machines sufficiently flexible to adapt to the environment and the different states of their human teammates, to anticipate the human teammates' capabilities and intentions, and to generalize from learned experiences to new situations.[40]

The U.S. defense budget directly mentions the development of "human/machine teaming" in allocating funds to autonomy research, and human-machine collaboration and teaming is a crosscutting theme across autonomy and AI research programs related to unmanned systems, information processing, decision support, targeting functions, and other areas. It includes initiatives such as:

- Air Force research on "Human Dynamics Evaluation" looks to "enable and enhance airman-machine teaming for distributed multi-domain operations."[41]

- Army basic research on "Human Engineering" explores novel forms of joint human-intelligent agent decision-making for "improved, emergent group performance," focusing on "deep learning approaches that function under conditions of limited, mismatched, or dynamic data." [42]

- Within the Army's applied NGCV program, the "Joint Human-Agent Teamwork" component "focuses on providing human intelligent agent teams that have the capability to perform as well as soldier teams, but with additional capabilities including: greater team resilience with robust and adaptive performance, faster dynamic human-agent team reconfiguration...faster and more informed team decision making, and reduced numbers of soldiers as well as risks to them."[43]

- Another applied NGCV research initiative looks to new "ML and reinforcement learning methods" to enable "joint human-intelligent agent decision making, optimizing the strengths of each in the decision process and creating an adaptive, agile team."[44]

- Within Army's applied research on "Soldier Lethality Technology", the "Soldier System Interfaces/Integration – Sensor Technology" project investigates the development of advanced user interfaces to "optimize human-robotic interaction during dismounted operations," and explores technologies and algorithms to "enable Squad and Platoon level autonomous reconnaissance using robotic systems to minimize the operator's dedicated control of the systems and reduce their cognitive burden."[45]

- DARPA applied research on "Artificial Intelligence and Human Machine Symbiosis" develops technologies that "enable machines to function not only as tools that facilitate human action but as trusted partners to human operators." The project is especially interested in "systems that can understand human speech and extract information contained in diverse media; answer questions, reach conclusions, and propose explanations; and learn, reason, and apply knowledge gained through experience to respond intelligently to new and unforeseen events."[46]

Some research efforts seek specifically to understand human interactions with autonomous systems and intelligent agents, for example:

- Air Force basic "Decision Making" research focuses in part on "development and testing of advanced representations and processes for higher-level artificial intelligence, trust between humans and autonomous agents, mixed human-machine decision making."[47]
- Air Force "Human Effectiveness Applied Research" works on "Human Trust and Interaction," seeking to "understand, quantify, and calibrate trust factors influencing airman interactions with autonomous systems that can be applied to air-man machine teaming in future weapon systems."[48]

These examples are only a fraction of the research and development efforts exploring human-machine collaboration and teaming. Yet they illustrate that humans remain at the epicenter of U.S. military investments in autonomy and AI. Effective human-machine collaboration is therefore key to harnessing the full promise of AI for strengthening military capabilities and furthering U.S. strategic objectives.

Trends and gaps

A wide range of U.S. military research efforts center on using autonomy and AI in unmanned vehicles and systems, information processing, decision support, targeting functions, and other areas. This research cuts across the different joint warfighting functions, including command and control, information, intelligence, fires, movement and maneuver, protection and sustainment—clearly demonstrating DOD's vision to leverage these emerging technologies for an enduring competitive advantage across the competition-conflict continuum.

In addition to the budgetary analysis, our team conducted a close reading and manual review of the different research programs classified as related to autonomy and/or AI. In our assessment, there are at least two areas where gaps in investment could have negative effects on U.S. military capabilities and strategic interests: gaps in research on trust in human-machine teams, and a lack of consistent focus on the security and safety of AI systems.

In the discussion below, we refrain from including counts and dollar amounts associated with research programs on human-machine teaming or those focused on AI safety and security. There is no straightforward way to conduct a keyword search scanning the descriptions of PEs, projects, and components in our dataset to identify these particular initiatives without either missing many relevant efforts or including unrelated ones.[49] That said, in the process of writing this report, we shared our analysis with issue area experts, including those within DOD, who largely agreed with our assessment. Moreover, these gaps are important to highlight because, if left unaddressed, they could hinder U.S. military innovation and the development and fielding of AI systems in operational environments.

Trust in human–machine collaboration and teaming

As previously noted, the U.S. military has a broad range of autonomy and AI-related research programs pairing humans with machines. Trust affects the willingness of humans to use autonomous and AI-enabled systems and accept their outcomes and recommendations. As such, trust is essential for effective human-machine collaboration.

While we can't precisely estimate what portion of autonomy and/or AI research efforts include a human-machine collaboration element, it is clearly a major research area for the U.S. military given that the U.S. defense budget directly mentions the development of "human/machine teaming" in allocating funds to autonomy research. Yet only 18 out of the 789 components classified as related to autonomy and 11 out of the 287 components classified as related to AI mention the word "trust."

This gap is concerning because current evidence supports both the "trust gap" argument—which suggests humans often do not trust machines to perform effectively—and the "automation bias" argument—which posits that humans trust machines too much, especially in complex and high-risk situations.[50] Both present unique risks and challenges for the application of human-machine teams in military settings.

For instance, research shows that while algorithms generally outperform humans in forecasting and prediction tasks, people nonetheless trust and prefer humans' forecasts to those made by algorithms.[51] Moreover, people are quick to lose confidence in algorithms when the technology makes mistakes.[52] Resolving this "trust gap" and ensuring that human team members fully understand what the system can and cannot do in a given environment may be necessary before the U.S. military can deploy some of the AI technologies it is currently researching.

At the same time, because military decision-making often occurs under time-sensitive and stressful conditions, there are serious concerns about the implications of automation bias. Indeed, investigations into the fratricides caused by the semi-autonomous Patriot surface-to-air system in 2003 showed that they were caused in part due to "unwarranted and uncritical trust in automation."[53]

As AI advances allow for higher levels of autonomy in unmanned systems, human-machine teams will become more commonplace and expand to additional areas and domains. Yet without a comprehensive and contextualized understanding of what drives trust in human-machine teams, the U.S. military may fail to effectively employ intelligent agents and systems as trusted partners to human operators. This obstacle, in turn, could stymie progress toward fielding AI as an enabler on future battlefields and impair the U.S. ability to achieve its broader strategic goals.

Trustworthy, robust, responsible, and secure AI

In our review of U.S. military S&T research on autonomy and AI, we identified several programs focused explicitly on increasing AI system robustness and resilience, strengthening security in the face of deceptive and adversarial attacks, and developing systems that behave reliably in operational settings. Yet of the hundreds of autonomy and AI research programs we reviewed, most don't mention safety and security attributes. Once again, while we refrain from quantifying this assessment through a keyword query, we believe it is worth sharing. In the short term, these gaps could hinder efforts to scale AI for decision support and undercut effective human-machine teaming; over the long term, they could also stall the development and fielding of AI systems in operational environments.

A closer look at the program descriptions shows that DARPA leads in research on making AI safe and secure. In basic research, DARPA programs include the "Foundational Artificial Intelligence Science," "Guaranteeing AI Robustness against Deception (GARD)," "Machine Common Sense (MSC)," "Mining and Understanding Software Enclaves (MUSE)," "Learning with Less Labels (LwLL)," "Lifelong Learning Machines (L2M)," and "Transparent Computing." DARPA's applied "Artificial Intelligence and Human–Machine Symbiosis" project includes initiatives such as "Accelerating Artificial Intelligence," "Assured Autonomy," and "Explainable Artificial Intelligence (XAI)."[54]

The Army also has a number of relevant initiatives. As part of its basic research portfolio, the "Army Collaborative Research and Tech Alliances" effort includes research on "AI-enabled cyber security that is robust to enemy deception," supporting "Army counter-AI against near-peer adversaries."[55]

Its applied research focuses on "Explainable Intelligence Underlying Efficient Integration of Cognitive Assist Agents," as well as on robust cybersecurity as part of "Defensive Cyber Operations."[56] Under the Next Generation Combat Vehicle PE, research on "Autonomous System Modeling and Simulations" seeks "robust autonomy algorithm development."[57]

These and other similar programs are a positive step toward ensuring the safety, security, and robustness of DOD AI systems. Such research efforts also reflect the DOD's commitment to complying with the five principles for ethical use of AI adopted in February 2020. They call for responsible, equitable, traceable, reliable, and governable AI for both combat and non-combat purposes.[58]Still, in our assessment, while there are several research programs explicitly dedicated to enhancing AI safety and security, most don't mention safety and security features such as robustness, resilience, reliability, trustworthiness, traceability, explainability, interpretability, or assurance.

Our analysis is limited by the research program descriptions in the justification books we examined. In addition to space constraints, program managers can only share so much about these research efforts in public documents. These limitations notwithstanding, when considered alongside the aforementioned gaps in research on trust in human-machine teaming, the implications of neglecting a consistent emphasis on safety and security in autonomy and AI research merit consideration.

- **Trustworthy AI systems are key to effective human-machine collaboration:** The U.S. military wants to employ machines as trusted partners to human operators. Research programs focused specifically on human-machine collaboration should therefore emphasize AI trustworthiness—including factors such as explainability, transparency, and reliability. Trustworthiness should be a crosscutting element across different autonomy and AI related research efforts that include human-machine teaming in some capacity.

    - Without emphasizing trustworthiness, AI applications for decision support could prove difficult to adopt and deploy. Decision support technologies integrate and fuse information from diverse sources, detect anomalies and changes, and provide recommendations for action. Human operators need to understand the logic behind outcomes to have confidence in AI solutions. If the commander does not understand how or why the system has arrived at a given conclusion, they may be reluctant to follow its recommendations.

- **Research and development programs related to AI in operational environments require an emphasis on robustness and reliability:** U.S. military investments in autonomy and AI related research—particularly at the applied and advanced levels—demonstrate an intent to develop AI for use in contested and hostile environments, across a range of missions and domains. Current AI systems, however, are largely unsuited for operational deployment; they are vulnerable to adversarial manipulation and attack, and cannot reliably handle uncertain and novel situations. As researchers work on these problems, emphasizing robustness and resilience early in the technology's lifecycle could help alleviate some of the challenges AI/ML systems will likely face during the testing, evaluation, verification, and validation stages.

As a whole, the budgetary assessment of U.S. military S&T priorities for autonomy and AI indicates progress toward leveraging these technologies to provide the U.S. military with the capabilities needed to deter adversaries, perform effectively on the battlefield, and collaborate with allies. But a closer reading of these research programs draws attention to gaps in our understanding of trust in human-machine collaboration and the limited attention paid to robustness, reliability, and interpretability of AI systems. These and other oversights could impede U.S. ability to integrate autonomy and AI into its military systems and operations. As a result, the capability enhancements and strategic advantages expected from AI—from stronger deterrence to battlefield effectiveness to dynamic alliances—may fail to materialize.

Conclusion

The U.S. military is investing in a broad range of research and development efforts applying autonomy and AI to unmanned systems, information processing, decision support, targeting functions, and other areas. These investments are often explicitly geared toward enhancing military capabilities such as speed, precision, coordination, reach, persistence, lethality, and endurance. On a strategic level, these technological advances and capability enhancements are ultimately meant to strengthen deterrence, ensure military effectiveness in future conflicts, and facilitate interoperability with U.S. allies.

While national security leaders identify artificial intelligence as a priority technology for defending the United States, our analysis shows it is difficult to accurately assess how much the U.S. military is investing in research related to autonomy and AI. Still, the data suggests that the share of the U.S. military science and technology program budget dedicated to AI research unrelated to autonomy and unmanned systems is relatively small. We also find gaps in investments in basic AI research, which raises concerns about U.S. military innovation in an increasingly competitive security environment.

The U.S. military plans to use AI in physical systems—in the air, on the ground, underwater, and in space—as well as to operate virtually in cyber operations and electronic warfare. Public discussions about this future of military AI tend to overemphasize the inevitability of fully autonomous weapon systems, swarms of miniature self-directed drones, and the transfer of battlefield decision-making from humans to machines. Our analysis, however, shows that the U.S. military is developing AI applications to augment rather than replace human intelligence. Indeed, human-machine collaboration is a prominent theme across the different research programs related to autonomy and AI. That said, gaps in research on trust in human-machine teams, and lack of consistent focus on security and safety of AI systems could all hinder U.S. military innovation and stymie the development and fielding of AI systems in operational settings.

Today’s investments will set the course for the future of AI in national security. By prioritizing reliable and secure AI systems, the United States will be able to ensure long-term military, technological, and strategic advantages.

Acknowledgments

In the process of research and writing, the authors benefited from the insights and feedback of many colleagues and experts. Thanks to Jason Matheny, Igor Mikolic-Torreira, Dewey Murdick, Helen Toner, Lynne Weil, Melissa Flagg, and Andrew Imbrie for their valuable comments. We are grateful to Danielle Tarraf of RAND and Martijn Rasser of the Center for a New American Security for reviewing the report and providing thorough and thoughtful comments that significantly strengthened our work. Many thanks also to Philippe Loustaunau for data support, Gina Acevedo and Tarun Krishnakumar for research support, Farhana Hossain for help with data visualization, and Matt Mahoney and Alexandra Vreeman for editorial support.

Document Identifier: doi: 10.51593/20200069

Appendix I: Definitions

This appendix compiles existing definitions of artificial intelligence, autonomy, autonomous weapons, and weapon systems across the Department of Defense that were encountered throughout the project.

Artificial Intelligence

Department of Defense Artificial Intelligence Strategy 2018

"AI refers to the ability of machines to perform tasks that normally require human intelligence—for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action—whether digitally or as the smart software behind autonomous physical systems."[59]

Defense Science Board

The DSB conducted a study on autonomy in which they explain "the primary intellectual foundation for autonomy stems from artificial intelligence (AI), the capability of computer systems to perform tasks that normally require human intelligence (e.g., perception, conversation, decision making)."[60]

Defense Innovation Board

The DIB published recommendations for AI principles in which the group defined AI as "a variety of information processing techniques and technologies used to perform a goal-oriented task and the means to reason in the pursuit of that task."[61]

DARPA

John Launchbury, former Director of DARPA's Information Innovation Office, defined AI as "a programmed ability to process information." He highlights three consecutive waves of AI: handcrafted knowledge (expert systems in the 1980s), statistical learning (machine learning and neural network models), and contextual adaptation (next evolution of AI). These three waves demonstrate AI progress along four dimensions: perception, learning, abstraction, and reasoning.[62]

FY2019 National Defense Authorization Act

1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.

2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.

3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.

4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.

5) An artificial system designed to act rationally, including and intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.[63]

National Security Commission for Artificial Intelligence

NSCAI uses the FY2019 NDAA definition of AI and further defines AI in its interim report as “the ability of a computer system to solve problems and to perform tasks that would otherwise require human intelligence.”[64]

Autonomy

DOD Community of Interest on Autonomy

"At the most basic level, autonomy draws on three broad, multidisciplinary technical fields: perception, cognition, and action."[65]

National Defense Industrial Association

At the NDIA 19th Annual Science & Engineering Technology Conference, the Air Force Autonomy Colonel Lead defined autonomy as “the computational capability for intelligent behavior that can perform complex missions in challenging environments with greatly reduced need for human intervention, while promoting effective man-machine interaction."[66]

DARPA

"Autonomy refers to a system's ability to accomplish goals independently, or with minimal supervision from human operators in environments that are complex and unpredictable."[67]

DEFENSE INNOVATION BOARD

The DIB did not explicitly define autonomy in its recent AI principles publication, but the group recognizes that AI is not equivalent to autonomy.[68]

DEFENSE SCIENCE BOARD

The DSB offers two definitions of autonomy in two separate reports. The first report, *Task Force Report on the Role of Autonomy in DoD Systems*, published July 2012, states that autonomy is a capability (or a set of capabilities) that enables a particular action of a system to be automatic or, within programmed boundaries, "self-governing."[69]

The second report, *Study on Autonomy*, published June 2016, asserts that autonomy "results from delegation of a decision to an authorized entity to take action within specific boundaries" and distinguishes between systems that are automated and autonomous. For a system to be autonomous, DSB states that it needs "to have the capability to independently compose and select among different courses of action to accomplish goals based on its knowledge and understanding of the world, itself, and the situation.[70]

STOCKHOLM INTERNATIONAL PEACE RESEARCH INSTITUTE

In November 2017, SIPRI published the report *Mapping the Development of Autonomy in Weapons Systems,* where it defines autonomy as "the ability of a machine to execute a task, or tasks, without human input, using interactions of computer programming with the environment." SIPRI went further to describe an autonomous system as "by extension, usually understood as a system—whether hardware or software—that, once activated, can perform some tasks or functions on its own."[71]

Weapon System

DOD DICTIONARY OF MILITARY AND ASSOCIATED TERMS ON WEAPON SYSTEM

"A combination of one of more weapons with all related equipment, materials, services, personnel, and means of delivery and deployment (if applicable) required for self-sufficiency (JP 3-0)."[72]

Autonomous Weapons

DEPARTMENT OF DEFENSE DIRECTIVE 3000.09

Lethal Autonomous Weapon Systems do not have a singularly accepted and explicit definition mainly because of the range of ways in which autonomous platforms are generally defined. Following are the definitions given by the 3000.09 Directive of the DOD:[73]

- ***Autonomous weapons systems***: "A weapon system(s) that, once activated, can select and engage targets without further intervention by a human operator. This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation."

- ***Human-supervised autonomous weapon system:*** "An autonomous weapon system that is designed to provide human operators with the ability to intervene and terminate engagements, including in the event of a weapon system failure, before unacceptable levels of damage occur."[74]

- ***Semi-autonomous system:*** "A weapon system that, once activated, is intended to only engage specific targets or specific target groups that have been selected by the human operator." The document specifies fire-and-forget homing munitions such as guided air-to-air missiles as examples.

"Semi-autonomous weapon systems that employ autonomy for engagement-related functions including, but not limited to, acquiring, tracking, and identifying potential targets; cueing potential targets to human operators; prioritizing selected targets; timing of when to fire; or providing terminal guidance to home in on selected targets, provided that human control is retained over the decision to select individual targets and specific target groups for engagement.

"'Fire and forget' or lock-on-after-launch homing munitions that rely on TTPs to maximize the probability that the only targets within the seeker's acquisition basket when the seeker activates are those individual targets or specific target groups that have been selected by a human operator."

Appendix II: Additional Tables and Figures

Figure 1A. U.S. Army S&T investments in autonomy and AI: component, PE, and maximum estimates (USD in millions, FY2018–FY2020)

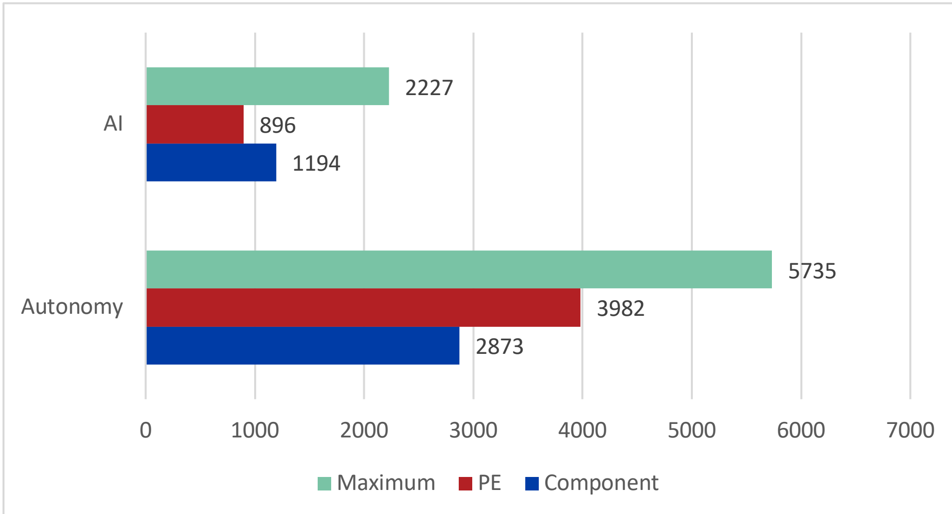

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Figure 1B. U.S. Navy S&T investments in autonomy and AI: component, PE, and maximum estimates (USD in millions, FY2018–FY2020)

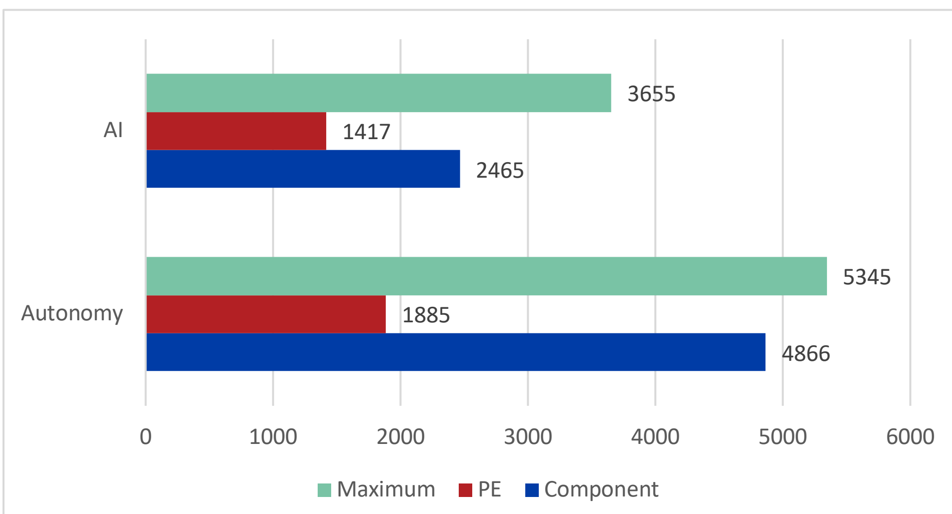

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Figure 1C. U.S. Air Force S&T investments in autonomy and AI: component, PE, and maximum estimates (USD in millions, FY2018–FY2020)

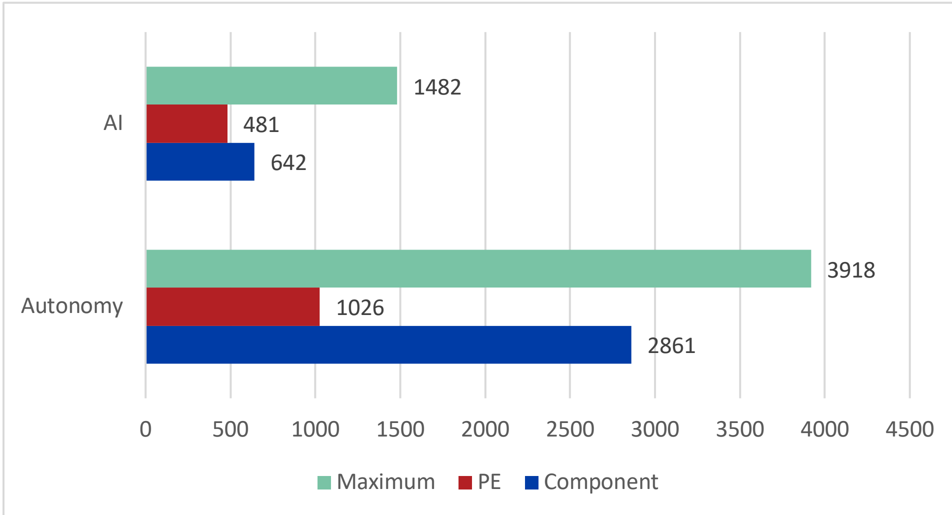

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Figure 1D. DARPA S&T investments in autonomy and AI: component, PE, and maximum estimates (USD in millions, FY2018–FY2020)

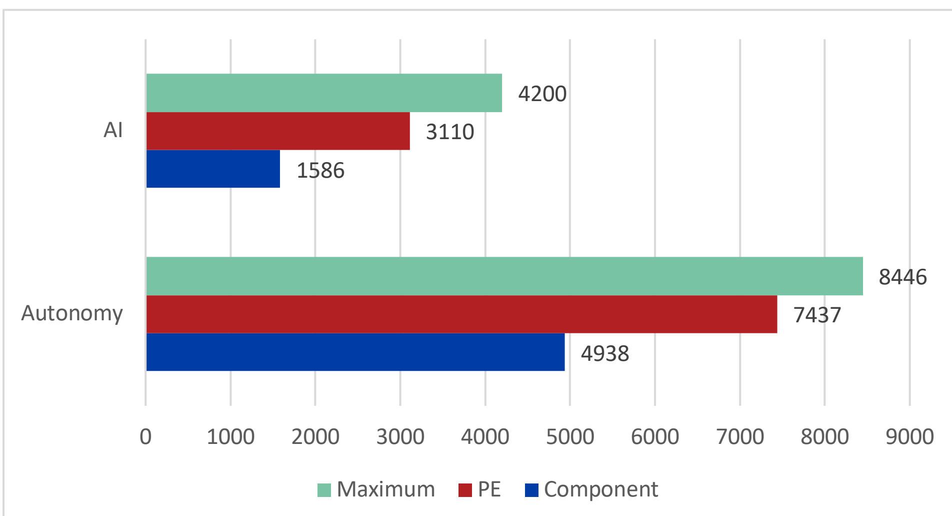

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Table 1A. Autonomy and AI-related PEs, projects, and components, counts by service

Autonomy			
	PE	Project	Component
Army	30	159	436
Navy	2	10	108
Air Force	4	21	106
DARPA	8	11	139
Total	44	201	789
AI			
	PE	Project	Component
Army	8	44	155
Navy	1	4	49
Air Force	1	7	33
DARPA	3	7	50
Total	13	62	287
Autonomy and AI			
	PE	Project	Component
Army	6	24	100
Navy	1	3	31
Air Force	0	4	33
DARPA	2	5	34
Total	9	36	198

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Table 2A. Autonomy and AI-related PEs, projects, and components, counts by research category

Autonomy			
	PE	Project	Component
Basic	2	23	142
Applied	24	97	383
Advanced	18	81	264
Total	44	201	789
AI			
	PE	Project	Component
Basic	3	14	77
Applied	8	30	152
Advanced	2	18	58
Total	13	62	287
Autonomy AND AI			
	PE	Project	Component
Basic	2	10	59
Applied	6	18	103
Advanced	1	8	36
Total	9	36	198

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Figure 2A. Assessing the overlap between autonomy and AI-related research, component counts by service and research category

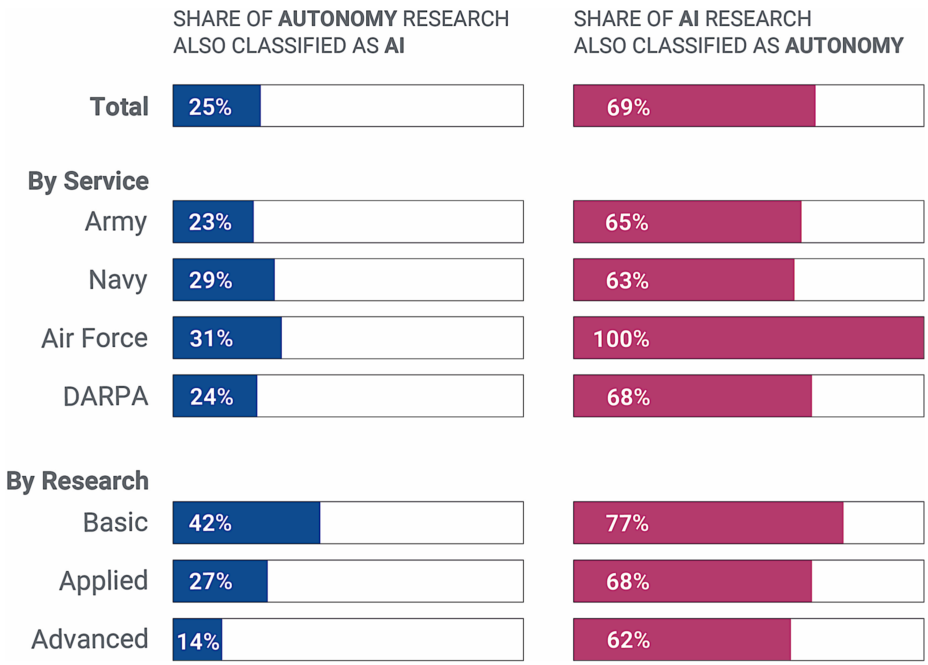

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Table 3A. U.S. Military investments in basic, applied, and advanced AI research: annual PE and maximum estimates (2018–2024, USD in millions)

	2018	2019	2020	2021	2022	2023	2024
				Basic Research			
PE	1,009	1,087	1,061	1,065	1,063	1,053	1,056
Max.	1,220	1,325	1,335	1,302	1,304	1,301	1,307
				Applied Research			
PE	815	851	880	909	916	975	996
Max.	1,704	1,893	1,990	1,387	1,387	1,456	1,509
				Advanced Research			
PE	51	44	107	130	136	146	145
Max.	649	635	814	545	550	618	626

Source: Department of Defense FY2020 Budget Estimates, RDT&E Justification Books of the U.S. Army, Navy, Air Force, and DARPA.

Endnotes

[1] Some of the recommendations in this report were also articulated in a report published jointly by the Bipartisan Policy Center and CSET. See Bipartisan Policy Center and the Center for Security and Emerging Technology, *Artificial Intelligence and National Security*, (Washington, DC: BPC, CSET, June 2020), https://bipartisanpolicy.org/wp-content/uploads/2020/07/BPC-Artificial-Intelligence-and-National-Security_Brief-Final-1.pdf.
[2] Congressional Research Service, *Federal Research and Development (R&D) Funding: FY2020* (Washington, DC: CRS updated March 2020), 3, https://fas.org/sgp/crs/misc/R45715.pdf.
[3] Office of the Under Secretary of Defense (Comptroller), *Defense Budget Overview: United States Department of Defense Fiscal Year 2021 Budget Request* (Washington, DC: Department of Defense, February 2020), 1–8, https://comptroller.defense.gov/Portals/45/Documents/defbudget/fy2021/fy2021_Budget_Request_Overview_Book.pdf; The FY2020 U.S. defense budget allocated $3.7 billion to unmanned and autonomous systems that will enhance "freedom of maneuver and lethality in contested environments through development of offensive-armed Unmanned Surface Vessel, Unmanned Undersea Vehicle, and Autonomous Logistics Platforms," and $927 million for AI and machine learning projects through the Joint Artificial Intelligence Center (JAIC) and Advanced Image Recognition (Project Maven).
[4] U.S. Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy: Harnessing AI to Advance Our Security and Prosperity* (Washington, DC: Department of Defense, 2018), https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF.
[5] National Security Commission on Artificial Intelligence, *Interim Report*, (Washington, DC: November 2019), https://www.nscai.gov/reports.
[6] Congressional Research Service, *Artificial Intelligence and National Security* (Washington, DC: CRS updated November 2019), https://fas.org/sgp/crs/natsec/R45178.pdf.
[7] Defense Science Board, "Summer Study on Autonomy," (Washington, DC: June 2016), 4, https://fas.org/irp/agency/dod/dsb/autonomy-ss.pdf.
[8] Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy* (Washington, DC: 2018), 5, https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF.
[9] Mary T. Tyszkiewicz and Stephen Daggett, *A Defense Budget Primer*, CRS Report No. RL30002 (Washington, DC: Congressional Research Service, 1998), https://fas.org/sgp/crs/natsec/RL30002.pdf.
[10] Donna Fossum, Lawrence S. Painter, Valerie L. Williams, Allison Yezril and Elaine M. Newton, "Government-Wide and DOD Definitions of R&D," in *Discovery and Innovation* (Santa Monica, CA: RAND, 2000), 615, https://www.rand.org/content/dam/rand/pubs/monograph_reports/MR1194/MR1194.appb.pdf.
[11] We developed the list of keywords by consulting numerous studies focused on military applications of autonomy and AI cited throughout this report. We also iterated between reading the program descriptions in the justification books and developing the list of

keywords to ensure we accounted for the different ways each service describes their research efforts.

[12] The AI keywords list was originally developed as part of a CSET research project on different definitions of AI in scientific literature. See James Dunham, Jennifer Melot, and Dewey Murdick, "Identifying the Development and Application of Artificial Intelligence in Scientific Text," February 17, 2020, https://arxiv.org/abs/2002.07143; Dewey Murdick, James Dunham, and Jennifer Melot, "AI Definitions Affect Policymaking," (Center for Security and Emerging Technology, June 2, 2020), https://cset.georgetown.edu/research/ai-definitions-affect-policymaking/.

[13] Fossum, et al., "Government-Wide and DOD Definitions of R&D," 615.

[14] Focusing on project level costs poses a similar problem because project budgets also encompass components not related to autonomy and AI.

[15] Bloomberg Government seems to use this approach in their analysis, referring to components as activities. Their analysis of the 2020 RDT&E budget identifies 222 AI related activities. Note, however, that while they use fewer keywords than included in our query, their analysis includes budget activities 4 though 7. Chris Cornillie, "Finding Artificial Intelligence Money in the Fiscal 2020 Budget," *Bloomberg Government*, March 28, 2019, https://about.bgov.com/news/finding-artificial-intelligence-money-fiscal-2020-budget/.

[16] It is possible that some of the discrepancy can be explained by the differences in how the services and DARPA allocate funds across the PE, project, and component categories.

[17] Ashwin Acharya and Zachary Arnold, "Chinese Public AI R&D Spending: Provisional Findings" (Center for Security and Emerging Technology, December 2019), https://cset.georgetown.edu/wp-content/uploads/Chinese-Public-AI-RD-Spending-Provisional-Findings-2.pdf.

[18] Defense Innovation Board, *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense* (Washington, DC: October 2018), 5, https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF. For instance, DOD Directive 3000.09 deals with autonomy in weapon systems, but does not address AI capabilities, in weapon systems or in any other context.

[19] United States Department of Defense, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Defense Advanced Research Projects Agency, Defense-Wide Justification Book Volume 1 of 5, Research, Development, Test & Evaluation, Defense-Wide* (Washington, DC: Department of Defense, 2019), vol 1–1, https://comptroller.defense.gov/Portals/45/Documents/defbudget/fy2020/budget_justification/pdfs/03_RDT_and_E/RDTE_Vol1_DARPA_MasterJustificationBook_PB_2020.pdf.

[20] United States Department of the Air Force, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Air Force, Justification Book Volume 1 of 3, Research, Development, Test & Evaluation, Air Force, Vol–I* (Washington, DC: Department of Defense, 2019), vol 1–17, https://www.saffm.hq.af.mil/Portals/84/documents/FY20/RDTE/FY20_PB_RDTE_Vol-I.PDF?ver=2019-03-18-153506-463.

[21] United States Department of the Navy, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Navy, Justification Book Volume 1 of 5, Research, Development, Test & Evaluation, Navy Budget Activities 1, 2, and 3* (Washington, DC: Department of Defense, 2019), vol 1–17, https://www.secnav.navy.mil/fmc/fmb/Documents/20pres/RDTEN_BA1-3_BOOK.pdf.

[22] Acharya and Arnold, "Chinese Public AI R&D Spending."

[23] A Congressional Research Service report on AI and national security, for instance, reviews AI applications for defense in areas such as intelligence, surveillance and reconnaissance, logistics, cyberspace operations, information operations, command and control, semiautonomous and autonomous vehicles, and lethal autonomous weapon systems. See Congressional Research Service, *Artificial Intelligence and National Security*.
[24] Danielle C. Tarraf, William Shelton, Edward Parker, Brien Alkire, Diana Gehlhaus Carew, Justin Grana, Alexis Levedahl, Jasmin Leveille, Jared Mondschein, James Ryseff, Ali Wyne, Dan Elinoff, Edward Geist, Benjamin N. Harris, Eric Hui, Cedric Kenney, Sydne Newberry, Chandler Sachs, Peter Schirmer, Daniel Schlang, Victoria M. Smith, Abbie Tingstad, Padmaja Vedula, and Kristin Warren, *The Department of Defense Posture for Artificial Intelligence: Assessment and Recommendations* (Santa Monica, CA: RAND, 2019), xii, http://www.rand.org/pubs/research_reports/RR4229.html.
[25] SIPRI's comprehensive study of autonomy in weapon systems finds that military systems today include multiple autonomous functions for mobility, targeting, intelligence, interoperability, and health management. See Vincent Boulanin and Maaike Verbruggen, "Mapping the Development of Autonomy in Weapon Systems" (Stockholm International Peace Research Institute, November 2017), 20, https://www.sipri.org/sites/default/files/2017-11/siprireport_mapping_the_development_of_autonomy_in_weapon_systems_1117_1.pdf.
[26] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2* (Washington, DC: Department of Defense, 2019), 130, https://www.asafm.army.mil/Portals/72/Documents/BudgetMaterial/2020/Base%20Budget/rdte/02%20RDTE%20-%20Vol%201%20-%20Budget%20Activity%202.pdf.
[27] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 3* (Washington, DC: Department of Defense, 2019), 99, https://www.asafm.army.mil/Portals/72/Documents/BudgetMaterial/2020/Base%20Budget/rdte/03%20RDTE%20-%20Vol%201%20-%20Budget%20Activity%202.pdf.
[28] United States Department of the Air Force, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Air Force, Justification Book Volume 1 of 3,* vol 1–105.
[29] United States Department of Defense, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Defense Advanced Research Projects Agency, Defense-Wide Justification Book Volume 1 of 5*, vol 1–61.
[30] Johan Schubert and Peter Svenmarck, *Artificial Intelligence for Decision Support in Command and Control Systems* (Swedish Defense Research Agency, November 2018), https://www.researchgate.net/profile/Johan_Schubert/publication/330638139_Artificial_Intelligence_for_Decision_Support_in_Command_and_Control_Systems/links/5c4b5f9a92851c22a3900b84/Artificial-Intelligence-for-Decision-Support-in-Command-and-Control-Systems.pdf; Karel van den Bosch and Adelbert Bronkhorst, "Human-AI Cooperation to Benefit Military Decision Making" (NATO, May 2018), https://www.researchgate.net/publication/325718292_Human-AI_Cooperation_to_Benefit_Military_Decision_Making.
[31] United States Department of the Navy, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Navy, Justification Book Volume 1 of 5*,1–273.

[32] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume 1, Budget Activity 1* (Washington, DC: Department of Defense, 2019), 1–8, https://www.asafm.army.mil/Portals/72/Documents/BudgetMaterial/2020/Base%20Budget/rdte/01%20RDTE%20-%20Vol%201%20-%20Budget%20Activity%201.pdf.
[33] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 178.
[34] Boulanin and Maaike, "Mapping the Development of Autonomy in Weapon Systems," 24, 26.
[35] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume 1, Budget Activity 1*, 13.
[36] United States Department of Defense, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Defense Advanced Research Projects Agency, Defense-Wide Justification Book Volume 1 of 5*, vol 1–236.
[37] United States Department of the Navy, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Navy, Justification Book Volume 1 of 5,* vol 1-207.
[38] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 174.
[39] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 3*, 443.
[40] Office of the Secretary of Defense, *Future Directions Workshop: Human Machine Teaming* (Arlington, VA: Department of Defense, 2020), https://basicresearch.defense.gov/Portals/61/Future%20Directions%20in%20Human%20Machine%20Teaming%20Workshop%20report%20%20%28for%20public%20release%29.pdf.
[41] United States Department of the Air Force, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Air Force, Justification Book Volume 1 of 3,* vol 1–67.
[42] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume 1, Budget Activity 1,* 44.
[43] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 128.
[44] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 133.
[45] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 81.
[46] United States Department of Defense, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Defense Advanced Research Projects Agency, Defense-Wide Justification Book Volume 1 of 5*, vol 1–82.
[47] United States Department of the Air Force, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Air Force, Justification Book Volume 1 of 3,* vol 1–11.

[48] United States Department of the Air Force, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Air Force, Justification Book Volume 1 of 3,* vol 1–67.
[49] Human-machine collaboration is a crosscutting theme across the different autonomy and AI related research programs, and developing a set of keywords that would capture the myriad of ways that humans can collaborate with intelligent technologies is likely to be extremely time consuming and would likely still undercount relevant efforts. Using the keywords query to identify research efforts related to AI safety and security, on the other hand, is likely to return many false positives given the frequent use of words such as "robust," "resilient," "security," "safety," and related terms in military parlance.
[50] Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey, "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General* 144, no. 1 (2015): 114–126.; Julia Macdonald and Jacquelyn Schneider, "Battlefield Responses to New Technologies: Views from the Ground on Unmanned Aircraft," *Security Studies* 28, no. 2 (February, 2019): 216–249.; John K. Hawley, "Patriot Wars: Automation and the Patriot Air and Missile Defense System" (Center for a New American Security, January 25, 2017), https://www.cnas.org/publications/reports/patriot-wars.
[51] Dietvorst, et al. "Algorithm Aversion,"; Michael C. Horowitz, Paul Scharre, and Alexander Velez-Green, "A Stable Nuclear Future? The Impact of Autonomous Systems and Artificial Intelligence," Arxiv, December 2019, https://arxiv.org/pdf/1912.05291.pdf.
[52] Julia L. Wright, Jessie Y.C. Chen, Shan G. Lakhmani, "Agent Transparency and Reliability in Human–Robot Interaction: The Influence on User Confidence and Perceived Reliability," *IEEE Transactions on Human-Machine Systems* 50, no. 3 (June 2020): 254–263, doi: 10.1109/THMS.2019.2925717.
[53] Mick Ryan, "Human-Machine Teaming for Future Ground Forces" (Washington, DC: Center for Strategic and Budgetary Assessments, 2018), https://csbaonline.org/uploads/documents/Human_Machine_Teaming_FinalFormat.pdf; Paul Scharre, *Army of None* (New York: W. W. Norton & Company, 2018), 324; Hawley, "Patriot Wars."
[54] United States Department of Defense, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Defense Advanced Research Projects Agency, Defense-Wide Justification Book Volume 1 of 5*.
[55] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 1*, 182.
[56] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2*, 506, 528.
[57] United States Department of the Army, *Fiscal Year (FY) 2020 Budget Estimates, March 2019, Army, Justification Book of Research, Development, Test & Evaluation, Army RDT&E – Volume I, Budget Activity 2,* 121.
[58] Defense Innovation Board, *AI Principles.*
[59] Department of Defense, *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*, 5.
[60] Defense Science Board, "Summer Study on Autonomy," 4.
[61] Defense Innovation Board, *AI Principles*, 5.
[62] John Launchbury, A DARPA Perspective on Artificial Intelligence, DARPA, https://www.darpa.mil/attachments/AIFull.pdf.

[63] John S. McCain National Defense Authorization Act for Fiscal Year 2019, H.R. 5515, 115th Congress, (2018).
[64] National Security Commission on Artificial Intelligence, *Interim Report*, 7.
[65] Office of Technical Intelligence, *Technical Assessment: Autonomy,* Department of Defense (Washington, DC: February 2015), 1, https://defenseinnovationmarketplace.dtic.mil/wp-content/uploads/2018/02/OTI_TechnicalAssessment-AutonomyPublicRelease_vF.pdf.
[66] Kris Kearns, *DOD Autonomy Roadmap*, Department of Defense (Austin, TX: March 21, 2018), 4, https://ndiastorage.blob.core.usgovcloudapi.net/ndia/2018/science/Kearns.pdf.
[67] Sandeep Neema, "Assured Autonomy," DARPA, https://www.darpa.mil/program/assured-autonomy.
[68] Defense Innovation Board, *AI Principles*, 5.
[69] Defense Science Board, "The Roles of Autonomy in DOD Systems," (Washington, DC: July 2012), 1, https://fas.org/irp/agency/dod/dsb/autonomy.pdf.
[70] Defense Science Board, "Summer Study on Autonomy," 4.
[71] Boulanin and Vergruggen, "Mapping the Development of Autonomy in Weapon Systems," 5.
[72] Department of Defense, *DOD Dictionary of Military and Associated Terms*, last updated January 2020, 230, https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/dictionary.pdf.
[73] Undersecretary for Defense (Policy), *Autonomy in Weapon Systems*, DOD 3000.09, November 21, 2012, https://www.hsdl.org/?view&did=726163.
[74] DODD 3000.09 specifically indicates that it does not cover "autonomous or semi-autonomous cyberspace systems for cyberspace operations; unarmed, unmanned platforms; unguided munitions; munitions manually guided by the operator (e.g., laser- or wire-guided munitions); mines; [or] unexploded explosive ordnance."