




NATIONAL
SECURITY
COMMISSION
ON ARTIFICIAL
INTELLIGENCE

Interim Report

NOVEMBER 2019



Commission Members

DR. ERIC SCHMIDT

Chairman

HON. ROBERT O. WORK

Vice Chairman

SAFRA CATZ

DR. STEVE CHIEN

HON. MIGNON CLYBURN

CHRISTOPHER DARBY

DR. KENNETH FORD

DR. JOSE-MARIE GRIFFITHS

DR. ERIC HORVITZ

ANDREW JASSY

GILMAN LOUIE

DR. WILLIAM MARK


DR. JASON MATHENY

HON. KATHARINA MCFARLAND

DR. ANDREW MOORE

Contents

I.	Message from the Chairman and Vice Chairman	1
II.	Preface.....	4
III.	The Challenge Before Us.....	6
	What Do We Mean by “AI”?	7
	Why Does AI Matter?	8
	How Could AI Advance National Security?	9
	What Threats Does AI Pose?	11
	A Period of Uncertainty and Debate	13
	Alternative Visions of the Future	14
	Agreeing on Basic Principles	15
	Trendlines of Concern.....	17
	The China Entanglement Challenge	18
	American Advantages.....	20
	The State of AI in Government	21
IV.	Lines of Effort for the U.S. Government.....	24
	1. Invest in AI Research and Development	24
	2. Apply AI to National Security Missions.....	29
	3. Train and Recruit AI Talent.....	35
	4. Protect and Build Upon U.S. Technology Advantages	40
	5. Marshal Global AI Cooperation	44
V.	Considerations on Ethical and Trustworthy AI.....	48
VI.	“Associated Technologies”	50
	<i>Appendix 1: Technical Discussion: What Is AI?</i>	53
	<i>Appendix 2: DoD-Tech Sector “Business Challenges”</i>	59
	<i>Appendix 3: AI Workforce Model</i>	61
	<i>Appendix 4: Organizations Consulted</i>	66
	<i>Appendix 5: Commission Staff and Advisors</i>	68





I. *Message from the Chairman and Vice Chairman*

We are pleased to provide Congress with the Interim Report of the National Security Commission on Artificial Intelligence. It represents the Commission’s initial assessment on artificial intelligence (AI) as it relates to national security, provides preliminary judgments regarding areas where the United States can do better, and suggests some interim actions the government could take now. Our full analysis and recommendations will be made in our final report.

We are heartened by the bipartisan support that the Commission is receiving from Congress. The White House has been generous with its time and insights. Departments, agencies, and the Intelligence Community have provided resources, support, and have answered our questions. Everyone has been forthright about the government’s shortcomings and earnest in their determination to “get AI right.” We have enjoyed equal support from leaders in academia, civil society organizations, and the private sector. They have explained their roles in the AI ecosystem, outlined their concerns, and highlighted opportunities for utilizing AI for national security purposes. The Commission is eager to hear from many more Americans and our allies and partners as it examines the most significant national security dimensions of AI.

We attribute the widespread support across America to the basic proposition that all of the commissioners hold true: AI is integral to the technological revolution that we are now experiencing. How the United States adopts AI will have profound ramifications for our immediate security, economic well-being, and position in the world. Developments in AI cannot be separated from the emerging strategic competition with China and developments in the broader geopolitical landscape. We are concerned that America’s role as the world’s leading innovator is threatened. We are concerned that strategic competitors and non-state actors will employ AI to threaten Americans, our allies, and our values. We know strategic competitors are investing in research and application. It is only reasonable to conclude that AI-enabled capabilities could be used to threaten our critical infrastructure, amplify disinformation campaigns, and wage war. China has deployed AI to advance an autocratic agenda and to commit human rights

violations, setting an example that other authoritarian regimes will be quick to adopt and that will be increasingly difficult to counteract.

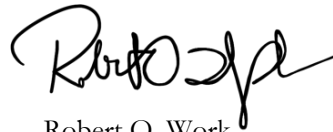
Given the robust and diverse views within the United States and the gravity of the challenge, we have developed seven consensus principles to guide our work and national discussion. First, global leadership in AI technology is a national security priority. The U.S. government retains a core responsibility to steer advancements in ways that protect the American people and ensure a robust basic research environment. Second, AI adoption for national security is an urgent imperative. We see no way to protect the American people, U.S. interests, and shape the development of international norms for using AI if the United States is not leading the way in application. Third, private sector leaders and government officials must build a shared sense of responsibility for the welfare and security of the American people. The government needs help from industry and academia to maximize the promise of AI and minimize the national security risks posed by AI. Fourth, people matter more than ever in the AI competition: we must cultivate homegrown AI talent and continue to attract the world's best minds. Fifth, actions taken to protect America's AI leadership from foreign threats must preserve principles of free inquiry, free enterprise, and the free flow of ideas. Sixth, at a basic level we see a convergence of interests and concerns between national security officials and those in the AI development and ethics community. Everyone wants safe, robust, and reliable AI systems; at the same time, today's technical limitations are widely recognized. Disagreements will persist, but we believe there is common ground that can serve as the basis for productive conversations. Seventh, any use of AI by the United States must have American values—including the rule of law—at its core.

In identifying areas of consensus, the Commission does not seek to downplay the contentiousness of many dimensions of the AI-national security nexus or minimize the complexity of policy choices. The deep interdependencies of the world's leading AI states present no easy answers for how best to further innovation, protect our security, and preserve our advantages. Finding concrete ways to protect U.S. companies and labs without undermining the principle of free inquiry is a hard problem. Ascertaining where the United States must reduce or constrain collaboration with China for human rights and national security purposes is a difficult challenge, given the deep interdependencies of the two AI communities within the larger worldwide community of AI researchers. The United States confronts hard choices between economic and security interests, between maintaining our openness and protecting our innovation economy from strategic competitors, and between commercial and national objectives, all while balancing short and long-term considerations.

The issues are too complex and vast for any part of government, society, or industry to address alone. Arriving at good solutions will require the work of the entire nation. Over the life of the Commission, we will continue to earnestly solicit diverse views as we seek answers to hard questions about the relationship between AI and national security. In our final report, we intend to make recommendations on how best to foster AI developments that will serve the interests of the American people, protect our national security, and uphold American values.



Eric Schmidt
Chairman



Robert O. Work
Vice Chairman

II. *Preface*

The Fiscal Year 2019 National Defense Authorization Act (NDAA) established the National Security Commission on Artificial Intelligence to “consider the methods and means necessary to advance the development of artificial intelligence, machine learning, and associated technologies by the United States to comprehensively address the national security and defense needs of the United States.”¹

Congress gave the Commission a broad mandate to examine artificial intelligence (AI) through the lens of national competitiveness, the means to sustain technological advantage, trends in international cooperation and competitiveness, ways to foster greater investment in basic and advanced research, workforce and training, potential risks of military use, ethical concerns, establishment of data standards, and the future evolution of AI.

The Commission will focus its inquiry on what it considers the most urgent challenges and the most transformative opportunities presented by AI for our national security. To date, our work has focused on:

- foreign threats to our national security in the current AI era;
- how AI can improve the government’s ability to defend the country, cooperate with allies, and preserve a favorable balance of military power in the world;
- the relationship between AI and economic competitiveness as a component of national security, including the strength of our scientific research community and our larger workforce; and
- ethical considerations in fielding AI systems for national security purposes.

This Interim Report fulfills Congress’s request for the Commission’s preliminary assessment of these challenges and opportunities. It is an attempt to inform policy and public debate about how developments in AI are related to wider national security trends. We are not yet in a position to make final recommendations, suggest major organizational changes, or propose specific investment priorities in rank order attached to dollar figures. We do believe, however, that laying out the basic fundamentals, presenting some consensus guiding principles, and offering initial preliminary judgments will contribute to public debate as the Commission moves toward its final report. This report identifies five fundamental lines of effort that are necessary to preserve U.S. advantages: Invest in AI Research and Development (R&D); Apply AI to National Security Missions; Train and Recruit AI Talent; Protect and Build Upon U.S. Technology Advantages; and Marshal Global AI Cooperation.

Between now and the publication of our final report, the Commission will pursue answers to hard problems, develop concrete recommendations on “methods and means” to integrate AI into national security missions, and make itself available to Congress and the executive branch to inform evidence-based decisions about resources, policy, and strategy.

Since the launch of the Commission in March 2019, we have engaged with Congress and the White House, and have enjoyed excellent cooperation from the Department of Defense, the Intelligence Community, and other parts of government. We have consulted with allied embassies in Washington. And our review has taken us to technology companies, non-profit groups, and universities across the country. We believe in soliciting advice from across all segments of society, industry, and government (please see [Appendix 4](#) for a list of many organizations consulted thus far). This is a national commission, and it requires seeking wisdom from across the nation. The Commission has heard from many Americans and is eager to hear from many more before it finalizes its recommendations.

III. *The Challenge Before Us*

The convergence of the artificial intelligence revolution and the reemergence of great power competition must focus the American mind. These two factors threaten the United States' role as the world's engine of innovation and American military superiority. If the United States fails to sustain its advantages, it will not be because it was caught by surprise. AI is not hidden in a top secret Manhattan Project. Tech luminaries and blue ribbon panels have sounded alarm bells on AI's peril and have championed its promise. Big countries and big tech companies with enormous quantities of data and computing power are leading the way, but the algorithms that fuel AI applications are publicly available. Open applications and development tools are diffusing at an accelerating pace. Many countries have published national AI plans. China, our most serious strategic competitor, has declared its intent to become the world leader in AI by 2030 as part of a broader strategy that will challenge America's military and economic position in Asia and beyond.²

The magnitude of technological change at a moment of strategic risk demands that our government and society find common purpose and face these challenges with the same imagination, decisive action, and national will summoned at other critical junctures in our history. The Commission believes Americans are up to the challenge.

We are optimistic that public officials will support AI investments to protect our national security and sustain our economic prosperity. We are confident that academia and private industry—especially universities and firms at the frontlines of AI research, development, and application—are willing

to reconceive their responsibilities for the health of our democracy and the security of our nation. And we see vitality in the American people's demand that their government pursue policies that maximize AI's potential, protect their privacy and civil liberties, and defend them from its malicious uses.

We are a bipartisan commission, with members from across the country. We are technology entrepreneurs, university leaders, and national security professionals. We are determined to build a common approach and inject the necessary energy for action in partnership with government officials, members of Congress, and the American people. Success requires that the Commission make the case for what it believes to be true: we

“We are in a strategic competition. AI will be at the center. The future of our national security and economy are at stake.”

are in a strategic competition. AI will be at the center. The future of our national security and economy are at stake.

WHAT DO WE MEAN BY “AI”?

AI is the ability of a computer system to solve problems and to perform tasks that would otherwise require human intelligence. AI technologies have evolved for many decades, including pattern recognition, machine learning, computer vision, natural language understanding, and speech recognition. These technologies are harnessed to enhance the abilities of both humans and machines, helping them to make decisions of higher quality and at greater speed. In a growing but still limited set of areas, machines can achieve human-like or better-than-human performance by analyzing large quantities of data, identifying patterns, and performing massive searches for useful answers, assessments, and recommendations. These systems are improving as the state-of-the-art shifts from expert systems based on explicit models to machine learning systems that can learn from experience and improve their performance, including those that can learn from sufficiently large and robust data sets. These are systems designed to solve tasks and achieve particular goals, with competencies that, in some respects, parallel the cognitive processes of humans: perceiving, reasoning, learning, communicating, deciding, and acting. [See [Appendix 1](#) for a more detailed discussion of the technical aspects of AI.]

The Foreseeable Future—Narrow AI and Human-Machine Teaming. Many AI researchers have been inspired by the possibility of developing systems that have the ability to perform all of the same intellectual tasks as humans. The term “artificial general intelligence” (AGI) is sometimes used to refer to this goal. Currently, people are the only example of general intelligence. Machines are far inferior at broadly learning, understanding and making inferences based on commonsense knowledge, generalizing from learnings on one task for handling another related problem, dealing with uncertainty, and leveraging concepts for logical and intuitive reasoning.³ When we might see the advent of AGI is widely debated. Rather than focusing on AGI in the near term, the Commission supports responsibly dealing with more “narrow” AI-enabled systems. Such technologies have already been developed and are being refined. They can be harnessed in powerful ways to augment human intelligence and allow new forms of human-machine collaboration and teaming. These developments and uses will assist humans in many different ways, including complementary autonomy on tasks that people find unsafe, undesirable, or unachievable alone. Responsibly utilizing today’s narrow AI applications as they apply to national security will pave the way and help prepare us for greater human-machine collaboration and machine-autonomy in the future.

Machine Learning: Over the past 25 years, AI has shifted from an era of reliance on explicit models created by experts to an era of statistical machine learning where engineers create statistical models with the capacity to be trained to perform within specific problem domains given exemplar data (e.g., images or sensor data) or simulated interactions (e.g., game playing). Today’s wave of rapid AI innovation is largely due to the emergence of a type of machine learning known as deep learning (DL). DL has proved to be an effective technique for image classification, object detection, speech recognition, and natural language processing, among other application areas. Learning from data, DL can solve problems never before solvable, or with a level of performance never before achievable. [See [Appendix 1](#) for further description of DL, its challenges, and machine learning alternatives.]

The Technical Core: AI is not a single piece of hardware or software, but rather, a constellation of technologies. One commissioner, Andrew Moore, describes AI as “a massive collection of interrelated technology blocks called the AI stack.”⁴ AI requires talent, data, hardware, algorithms, applications, and integration. We regard talent as the most valuable resource because it drives the creation and management of all of the other components. Data is critical for most AI systems.⁵ Labeled and curated data enables much of current machine learning used to create new applications and improve the performance of existing AI applications. The underlying hardware provides the computing power to analyze ever-growing data pools and run applications. Algorithms are the mathematical operations that tell the system how to navigate the data to provide answers in response to specific questions. An application makes the answers useful for specific tasks. Integration is critical to fielding a successful end-to-end AI system. This requires significant engineering talent and investment to integrate existing data flows, decision pipelines, legacy equipment, testing designs, etc. This task of integration can be daunting and historically has been underestimated.⁶

AI Ecosystem: Even with a strong technical core, efficient and effective development, acquisition, and deployment of AI rarely happens in isolation. It is important to have an ecosystem of support in place including laws, funding, institutions, policies, talent, intellectual property protection, supply chains, and counter-AI defenses.⁷

WHY DOES AI MATTER?

We have entered the age of AI. We take for granted the many ways it is changing our daily lives: helping us to navigate rush hour traffic; selecting just the right film on family movie night; or notifying us of fraudulent activity on our credit cards. AI will continue to transform society and the economy, although skepticism about AI’s potential is not

without precedent. Scientists recall two previous eras of AI promise when disappointing results led to “AI winters.” During these winters, research funding dried up and attention turned to other promising technologies. Today’s circumstances are different. Many of the techniques, algorithms, and theories developed in the past are now paying dividends because of breakthroughs in computational power, cloud computing, the availability of massive amounts of training data, improvements in machine learning algorithms, and mobile connectivity.⁸

The development of AI will shape the future of power. The nation with the most resilient and productive economic base will be best positioned to seize the mantle of world leadership. That base increasingly depends on the strength of the innovation economy, which in turn will depend on AI. AI will drive waves of advancement in commerce, transportation, health, education, financial markets, government, and national defense.

AI is “dual use” technology—usable for military and civilian purposes—but that simple dichotomy fails to capture AI’s ubiquity. AI is really a general purpose technology. There are no easy answers to the question of what to do about a general purpose technology that gives rise to outcomes that range from beneficial and benign to, potentially, existentially threatening.⁹ Facial recognition software tags friends’ photos posted on social media. In societies with strong commitments to civil liberties, it can make us safer. Yet it can also be employed on an industrial scale as a tool of repression and authoritarian control, as we have seen in China and elsewhere. Deepfakes—computer generated video or audio so sophisticated that it is indistinguishable from reality—produce harmless entertainment (like apps to superimpose your head on the bodies of a movie character), but could also be used to slander an individual or interfere in our political process.

HOW COULD AI ADVANCE NATIONAL SECURITY?

Adopting AI-enabled systems responsibly, rapidly, and at scale will allow national security organizations to understand and execute their missions faster. AI is not a panacea, but it will change the way we defend America; how we deter adversaries; how intelligence agencies make sense of the world; and how we fight.

AI Will Change How We Defend America. AI-enabled tools and systems can help protect our borders, detect and combat malicious cyber operations, safeguard our critical infrastructure, and respond effectively to natural disasters. The capability to process large amounts of data enables real-time risk assessment and response, and provides tailored situational awareness to first responders. Such tools can find the needle in the haystack, identifying anomalies to inform counterterrorism and counterintelligence efforts.¹⁰ AI-

equipped sensors at ports of entry will better prevent the illicit movement of weapons, contraband, and people. Predictive models and human-machine teams promise more rapid and efficient responses to threats and emergencies.

AI Will Change How Intelligence Agencies Make Sense of the World: In a world of widespread sensors, AI algorithms can sift through vast amounts of data to find patterns, detect threats, and identify correlations. AI tools can make satellite imagery, communications signals, economic indicators, social media data, and other large sources of information more intelligible. AI-enabled analysis can provide faster and more precise situational awareness that supports higher quality decision-making. The combination of AI capabilities and these new information sources also means that intelligence is now commercialized. Intelligence agencies are accustomed to operating in the shadows, but they are no longer alone. Private sector technologies give all types of actors access to what were once only government capabilities. Our national security agencies will need to understand the motivations and technical capabilities of adversaries—states, terrorists, and criminals—who may use AI to spy on us or do us harm.¹¹

AI Will Change How We Fight: On future battlefields, the military could use AI-enabled machines, systems, and weapons to understand the battlespace more quickly; develop a common joint operating picture more rapidly; make relevant decisions faster; mount more complex multi-domain operations in contested environments; put fewer U.S. service members at risk; and protect innocent lives and reduce collateral damage. AI will foster a new generation of semi-autonomous and autonomous combat systems and operations. Autonomous capabilities can be useful for a wide array of applications, including for predictive analysis, decision support systems, unmanned platforms, robotics, and weapons (both cyber and physical).¹²

The long-term strategic implications of adopting AI technologies for military purposes may be even greater than the impact on any specific military task. Because the integration of AI-enabled technology throughout military systems and operations will increase the accuracy and speed of perceiving, understanding, deciding, and acting well beyond the capability of human cognition alone, some believe AI will usher in a new era of “algorithmic warfare.”¹³ Algorithmic warfare will pit algorithms against algorithms in a contest dominated more by the speed and accuracy of knowledge and action than traditional factors such as force size, levels of armament, or the range of weapon systems. Battlefield advantage will shift to those with superior data, connectivity, compute power, algorithms, and overall system security. Reaching such a future will require the development of new operational concepts, organizational constructs, and decision-makers at all levels trained to understand AI and its associated technologies.

The Commission recognizes that lethal autonomous weapons systems (LAWS) represent an important dimension of public discourse about AI and national security. The Commission endeavors to understand different perspectives on LAWS and hear from different sides of the issue before attempting to reach consensus judgments.

WHAT THREATS DOES AI POSE?

One of the Commission's core responsibilities is to consider how an AI-enabled future could threaten U.S. security and interests. AI is intensifying and accelerating existing threats, while also creating novel threats. These could emerge from two vectors: what an adversary could do with AI, and what consequences an AI system could have if employed without safeguards, irrespective of who wields the technology.

- **Erosion of U.S. military advantage.** Strategic competitors, led by China and Russia, want to use AI-enabled autonomous systems in their military strategies, operations, and capabilities to undermine U.S. military superiority and conventional deterrence. Should they outpace the United States in these efforts, a basis for 70 years of strong alliances that forestalled great-power war and expanded global prosperity and freedom could be at risk.¹⁴
- **Strategic stability at risk.** Global stability and nuclear deterrence could be undermined if AI-enabled systems enable the tracking and targeting of previously invulnerable military assets. This dynamic could push states to adopt more aggressive force postures that upset existing bilateral, regional, or global stability, lead to the rapid escalation of conflict, or even increase incentives for a first strike.¹⁵
- **The diffusion of AI capabilities.** The likelihood of reckless and unethical uses of AI-enabled technologies by rogue states or non-state actors is increasing as AI applications become more readily available. Many of the algorithms and applications available in the public domain will have beneficial uses, but may also be utilized for malign ends. Criminals, terrorists, and lone wolves already empowered in the cyber era may be able to reach farther, faster, and with less attribution into our financial system, infrastructure, and personal lives.¹⁶
- **Disinformation and the threat to our democratic system.** AI will accelerate the already serious threat of cyber-enabled disinformation campaigns, including false-flag efforts. It will enable deepfakes—including live action computer-generated false realities—that can be distributed on a massive scale, with content

that exploits individual biases and preferences. Natural language processing can be leveraged to develop emotionally sophisticated messages micro-targeted to groups segmented by demographics and ideology. This could threaten the integrity of our public discourse and elections by leading citizens to make choices on the basis of falsehoods, and exacerbate existing divisions in society.¹⁷

- **Erosion of individual privacy and civil liberties.** New AI tools present states with greater capabilities to monitor and track their citizens or those of other states. While citizens' data can be used for lawful and legitimate purposes, the proliferation of new data sources—such as those generated through smart cities or smart policing—increases the risk of human rights abuses or violation of individual privacy. While China's use of AI surveillance tools has been well-documented, at least 74 other countries are also engaging in AI-powered surveillance, including many liberal democracies.¹⁸
- **Accelerated cyber attacks.** AI will advance traditional cyber capabilities that can move through multiple systems at superhuman speed. Intelligent malware will autonomously find and exploit system weaknesses, play offense and defense at the same time, and smartly target specific systems. AI also expands the threat vectors vulnerable to cyberattacks. As everyday systems from cars to critical infrastructure begin to rely on software, new defenses, likely relying on AI-enabled autonomous capabilities, will be needed to ensure these systems are reliable and secure.¹⁹
- **New techniques bring new vulnerabilities.** AI systems will be under constant stress from attackers using “counter-AI” techniques to pollute data, trick the machine, or reverse engineer information by gaining access to the algorithm. AI systems could be rendered ineffective if security is not built into AI systems from the beginning. Without trust and reliability, users will lose faith in AI's utility. Moreover, if AI can help us to understand adversaries, it also can help adversaries understand us. With AI tools to analyze individuals' digital or biometric data, adversaries could selectively target Americans for manipulation.²⁰
- **The danger of accidents.** Emerging technologies can generate safety and reliability risks. Some nations or actors may lower their safety and reliability standards and adopt AI-enabled technologies before they are ready, increasing the potential for mistakes, misperception, and unintended consequences. The chance of such accidents could be further compounded by AI's brittleness in complex environments that do not resemble a stable test environment. In addition, overreliance on AI systems could lead to preventable mistakes, and interactions

between multiple AI systems could result in unpredictable and catastrophic outcomes.²¹

A PERIOD OF UNCERTAINTY AND DEBATE

Given both the promise and peril of AI, and the fluid nature of technological development, we are in a period of uncertainty. The scope for positive action is wide, but there is a high degree of concern about the consequences of using AI-enabled technology, particularly for national security purposes. Over the course of the Commission's work we have heard how the private sector, non-governmental organizations, universities, and the government are wrestling with the current and potential effects of AI on society, the economy, and war. For example:

- Many defense experts urge the United States to move more quickly to field AI-enabled weapons systems and other capabilities, once the technology is ready, in light of growing international threats and the possibility that AI-enabled systems could save American lives and reduce civilian casualties. Conversely, some technologists and ethicists urge the United States to slow down or forswear the adoption of AI for military purposes, citing everything from a catastrophic accident, to crisis instability, to the immoral weaponization of AI.
- Some tech workers, concerned that their efforts will be used for military purposes, have called on their employers to limit the scope of their business with the Department of Defense (DoD) or avoid it altogether, while other companies have actively embraced the importance of working with the U.S. government for national security and defense purposes.
- Tech firms are establishing ethics committees to develop guidelines for fielding AI, protecting privacy, sharing data, and weighing whether or not to do business with the U.S. government or with countries that use AI to oppress their citizens. Uncertainty remains about how they will change company policy and behavior.
- Companies and universities stress the economic and intellectual importance of collaborative work with researchers from around the world, including from some countries that pose a strategic threat to the United States. This position stresses that the benefits of the overall accelerated advancement of scientific discovery, access to global talent, and economic gains may outweigh national security risks. National security and law enforcement officials stress the evidence that competitors, led by China, are extracting AI knowledge and technology from the

United States through licit and illicit means at a scope and scale that will undermine U.S. strategic advantages.

ALTERNATIVE VISIONS OF THE FUTURE

In light of these important differences, and in thinking about an uncertain future, the Commission must develop recommendations that maximize the promise of AI while reducing the risk of AI leading us down a perilous path. The Commission's attempts to predict AI's impact on national security is like Americans in the late 19th century pondering the impact of electricity on war and society. To move forward with purpose, we need a vision of the AI-empowered future that we seek to achieve, and envision AI futures we want to avoid for America and the world. A clarifying vision can help guide us through a complicated future.

- We envision a world in which AI is used to prolong and enrich lives and make people smarter by giving them the information they need when they need it. We hope AI can free people from dangerous and monotonous tasks so that more people can pursue meaningful and creative work. It will be a world where data and privacy remain protected from abuse and individual rights are preserved. AI-empowered systems will be used to diagnose disease and improve our health, forecast and improve our responses to natural disasters, and fuel scientific discoveries across all fields. In contrast to the fears of today, AI technologies can help us live in a world where individuals live freer from government coercion, state sovereignty is respected, and both individuals and states can pursue a more open exchange of ideas and goods. It will be a world where AI systems are used consistent with, and in service of, core values Americans hold dear and the rights enshrined in our founding documents.
- By contrast, we must avoid a future where AI contributes to a world of greater centralized control; empowers authoritarianism; is utilized as an instrument to repress dissent and impose conformity; destroys truth and trust within societies and between states; and is employed in reckless, irresponsible, and unethical ways.²² The careless weaponization of AI could be destabilizing, and the deliberate misuse of AI could do great harm. This may be of greatest risk with AI applications that are rushed into use without proper safeguards, sufficient testing, and consideration of ethics. We can glimpse that future by understanding what our competitors have already done and announced that they intend to do. China is using AI to build a dystopian surveillance state, and aspires to create social credit systems that assign people to “blacklists” based on who they

communicate with, where they travel, what they buy, and how they use their mobile phones.²³ It is leveraging facial recognition technology to identify and repress its minority populations.²⁴ Russia has already fielded armed robotic vehicles with autonomous features on the battlefield in support of a brutal dictator without evident regard for ethical considerations.²⁵ It will almost certainly use AI to accelerate its efforts to violate the sovereignty of other states using hybrid warfare.

AGREEING ON BASIC PRINCIPLES

In order to guide us toward a better future, the Commission hopes to begin focusing the current broad and passionate AI discussion by building consensus around seven basic principles about the relationship between AI and national security:

- **Global leadership in AI technology is a national security priority.** Given the centrality of AI to the future of our economy, society, and security, the U.S. government must pursue an investment strategy that extends America's technological edge. Global leadership gives our defense and security agencies access to the best technology, and puts the United States in the best position to secure that technology against vulnerabilities and develop international norms and standards for responsible use. While American companies play a significant role in advancing AI research and development, the government retains a core responsibility to steer advancements in ways that protect the American people and foster a robust basic research environment.
- **Adopting AI for defense and security purposes is an urgent national imperative.** Accelerating applications of AI to national security missions is an intelligence, warfighting, and organizational necessity. Our service members and officials must have access to the most advanced AI technologies to protect the American people, our allies, and our interests. The Commission is not glorifying the prospect of AI-enabled warfare. But new technology is almost always employed for the pursuit of power. In light of the choices being made by our strategic competitors, the United States must also examine AI through a military lens, including concepts for AI-enabled autonomous operations.
- **Private sector leaders and government officials must build a shared sense of responsibility for the welfare and security of the American people.** American companies are at the forefront of AI developments. This has profound national security consequences. Their investments dwarf federal R&D; they generate many of

the major breakthroughs; and they are on the frontlines of defending against cyber threats and malicious uses of AI applications. Industry must help government discern trends, act against foreign threats, and identify experts willing to help. The government must strengthen industry by articulating clear standards and policies for responsible use, rebuilding trust through greater transparency, and offering a vision of a shared purpose. To realize AI's potential, we must forge a common commitment to protecting our values, free market principles, and security.

- **People are still essential.** Talent remains the most important driver of progress in all facets of AI. We must prioritize cultivating homegrown talent by making long-term investments in STEM education. In the near term, high-skilled immigration is important for rapidly growing America's talent pool. One of America's advantages is the fact that its universities, companies, and innovation culture are magnets for the world's best AI talent. We need to encourage that talent to come, contribute, and stay. Within government, recruiting, training, and retaining AI-talent will be essential to maximize AI's potential.
- **The power of free inquiry must be preserved.** The open and collaborative AI innovation environment rests on the principles of free inquiry, free enterprise, and the free flow of ideas. We know other states are exploiting our open society. We must protect our intellectual property and sensitive technology. We must also ensure that American technology and innovation is not exploited to advance adversaries' militaries or undertake human rights abuses. However, any U.S. restrictions or controls must be narrowly tailored, linked to specific threats, and employed in ways that promote academic and commercial leadership in AI. Policies must fuel, not stifle, the innovation culture and values at the heart of our national power. Achieving protection and adequate safeguards, while maintaining openness poses a significant challenge because of the tension between the objectives. There are no easy answers.
- **Ethics and strategic necessity are compatible with one another.** Defense and national security agencies must develop and deploy AI in a responsible, trusted, and ethical manner to sustain public support, maximize operational effectiveness, maintain the integrity of the profession of arms, and strengthen international alliances. At a basic level we see a convergence of interests between national security officials and many in the AI development and ethics community. Everyone desires safe, robust, and reliable AI systems free of unwanted bias, and recognizes today's technical limitations. Everyone wants to establish thresholds for testing and deploying AI systems worthy of human trust and to ensure that humans remain responsible for the outcomes of their use. Some disagreements will remain, but the Commission is concerned that debate will paralyze AI development. Seen through the lens of

national security concerns, inaction on AI development raises as many ethical challenges as AI deployment. There is an ethical imperative to accelerate the fielding of safe, reliable, and secure AI systems that can be demonstrated to protect the American people, minimize operational dangers to U.S. service members, and make warfare more discriminating, which could reduce civilian casualties.

- **The American way of AI must reflect American values—including having the rule of law at its core.** For federal law enforcement agencies conducting national security investigations in the United States, that means using AI in ways that are consistent with constitutional principles of due process, individual privacy, equal protection, and non-discrimination. For American diplomacy, that means standing firm against uses of AI by authoritarian governments to repress individual freedom or violate the human rights of their citizens. And for the U.S. military, that means finding ways for AI to enhance its ability to uphold the laws of war and ensuring that current frameworks adequately cover AI.

TRENDLINES OF CONCERN

As the Commission considers alternate futures and affirms the imperative of sustaining U.S. leadership in AI based on consensus principles, it is concerned that trends in AI across economic and security fronts could lead to an AI future disadvantageous to American interests. Consider:

- *Research and Development:* China has overseen a 30 times increase in its overall R&D funding from 1991 to 2015, and is projected to surpass the United States in absolute R&D spending within 10 years.²⁶ U.S. federal investment in AI R&D has increased only marginally, as we discuss in greater detail below. Incrementalism will not assure U.S. leadership. America’s leadership in AI research—measured by indicators such as academic publications—is also shrinking. For example, one study found that Chinese researchers are “poised to overtake [their American counterparts] . . . in the most-cited 10% of papers next year, and in the 1% of most-cited papers by 2025.”²⁷
- *Commercial Competition:* Chinese tech firms have reached enormous scale and are poised to become leaders in applied AI, excelling in numerous commercial AI applications, including in healthcare, education, and e-commerce.²⁸ Some of these applications may pose national security risks.²⁹ China’s new “national team” of leading Chinese tech firms (including Baidu, Alibaba, Tencent, iFlytek, and SenseTime) is being harnessed to promote national objectives in AI, including by

supporting national laboratories working on deep learning, brain-inspired intelligence, and virtual/augmented reality.³⁰ The global reach and sophistication of these companies may soon eclipse American counterparts, giving Chinese firms access to the data, resources, and market power required to lead in AI.

- *Military-Civil Fusion:* China is intensifying efforts to exploit civilian and commercial developments in AI and leveraging a growing number of companies to advance Party-state and military purposes. The Chinese Communist Party’s concept of “military-civil fusion” has been elevated in national strategies and advanced through a range of initiatives. The distinction between civilian and military-relevant AI R&D is being eroded.³¹
- *Military Modernization:* China and Russia each have established research and development institutes to advance their military applications of AI, akin to the Defense Advanced Research Projects Agency (DARPA).³² Chinese researchers are developing military applications of AI technologies—including for swarming, decision support, and information operations—while the Chinese defense industry is pursuing the development of increasingly autonomous weapons systems. China is pursuing a process of “intelligentization” as a new imperative in military modernization.³³
- *Global Talent:* The United States is facing new competition for global STEM talent, especially in AI where there is a critical shortage of expertise.³⁴ China is undertaking an active effort to recruit global AI talent and persuade Chinese nationals working abroad to return to China.³⁵ Other countries are introducing favorable immigration and work policies to attract AI talent. We are beginning to see troublesome signs that America’s ability to attract and keep the top global talent may be weakening.³⁶

Without a reversal of current trends, in the coming decade the United States could lose its status as the primary base for global AI research, development, and application. If technological advances and AI adoption elsewhere outpace those in American firms and in the U.S. government, the resulting disadvantage to the United States could endanger U.S. national security and global stability.³⁷

THE CHINA ENTANGLEMENT CHALLENGE

China represents the most complex strategic challenge confronting the United States because of the many co-dependencies and entanglements between the two competitors.

Given the current trade tension, on top of a broader and growing global competition, many wonder if the United States should disentangle its economy and research network from China—including in the deeply interconnected field of AI.

China has long benefited from the open nature of the U.S. academic and commercial ecosystem to build its AI capacity, and Chinese nationals have made and continue to make valuable contributions to AI R&D in the United States. U.S. universities have trained many of China's premier AI researchers. Many American universities and technology companies depend on Chinese nationals for research and technical expertise.³⁸ U.S. and Chinese AI firms have partnered to do business in China, and Chinese companies have set up AI research labs and business subsidiaries in the United States. U.S. venture capitalists fund AI startups in China, while Chinese sovereign wealth funds, regional and local governments, universities, and companies are investing in and acquiring U.S. technology companies.

While AI-specific data is difficult to obtain, it is clear that a subset of Chinese nationals have engaged in intellectual property theft as well as state-directed espionage against the U.S. science and technology sectors. FBI Director Christopher Wray has stated that the bureau has “economic espionage investigations that almost invariably lead back to China in nearly all of our fifty-six field offices, and they span just about every industry or sector.”³⁹ According to the Department of Justice, “from 2011-2018, more than 90 percent of the Department's cases alleging economic espionage by or to benefit a state involve China, and more than two-thirds of the Department's theft of trade secrets cases have had a nexus to China.”⁴⁰ The threat extends beyond traditional espionage. Some researchers from China at U.S. universities have established parallel labs in China where they are commercializing technologies developed through U.S. research.⁴¹ Some researchers have taken advantage of the peer-review process for grant applications to illegally share or disclose information about U.S. research projects.⁴² China sends military scientists abroad to be trained in U.S. universities and to conduct research in AI and related technologies.⁴³

The Commission believes the United States must act to protect its interests from China's state-directed espionage, and protect against China's concerted efforts to steal or extract AI knowledge from private and public institutions. At the same time, in preliminary discussions with U.S. industry and academia, the Commission has heard about the commercial and research benefits of collaboration. U.S. industry and academic leaders have cautioned that the deep human, hardware, supply chain, investment, and corporate connections between the United States and China in AI cannot be unwound without economic costs and unintended consequences for the U.S. economy and U.S. research environment.⁴⁴ From this perspective, U.S. universities, labs, and companies could lose

access to valuable markets, an important talent pool, and the research now being generated by labs in China. Over the long term, the United States might be deprived of insights into Chinese advances and forgo the benefits of collaboration. Attempting to prevent American know-how or hardware from making its way to China could accelerate China's indigenous development or enable other nations' companies to profit from China's talent pool and the China market at the expense of U.S. companies.

The choice need not be a binary one between cooperating and disentangling. The Commission is seeking to better understand the specific actions that will balance competing interests and chart a sensible path forward for preserving beneficial elements of cooperation while establishing defenses against activities that run counter to American interests. The challenge the United States faces is how to recalibrate elements of the U.S.-China tech relationship to be more conducive to American interests and preserve U.S. advantages, taking into account realistic assessments of Chinese state-directed behavior, the need to mitigate intellectual property theft, and the importance of preventing the proliferation of technology used for human rights abuses.

AMERICAN ADVANTAGES

The Commission will present a more complete assessment of relative U.S. advantages and weaknesses in the final report. In the meantime, we must not lose sight of America's existing advantages in commercial and academic AI, and its historic advantages in nurturing a decentralized innovation ecosystem, drawing on the world's talent, and providing a system of government where free inquiry and risk-taking entrepreneurship are rewarded. U.S. universities remain the top centers for AI research.⁴⁵ The United States continues to attract, train, and retain the world's best for its companies and labs: around 80 percent of international computer science PhDs that are trained in the United States, including those from China, stay in the country after graduating.⁴⁶ American companies remain world leaders in AI research and some areas of application. Our market-based economy and low regulation has created three-quarters of the world's top 100 AI startups.⁴⁷ In all, we are home to more than 2,000 AI startups, twice that of our nearest competitor, and roughly half of the AI unicorns—private companies valued at more than a billion dollars.⁴⁸ When this assessment is broadened to include the combined assets of the United States and its allies, versus those of China and its partners, U.S. advantages are even more pronounced.

A successful American national security strategy for AI must leverage America's comparative advantages. The United States should acknowledge, but not overstate or seek to replicate, potential authoritarian advantages in AI. Authoritarian regimes can

amass and centralize data with little regard to privacy protections and compel ostensibly private companies to act on their behalf. Access to enormous quantities of data can provide benefits in using AI for specific fields like genomics. Authoritarian regimes have a greater propensity to issue centralized AI development plans backed by government funds and designate companies as “national champions.” They are more risk tolerant in fielding AI-systems quickly without the same ethical and legal safeguards that constrain democracies. However, on the technical side, the advantages of access to a larger data pool may be overstated. In the future, data will have diminishing returns as algorithms improve and it is replaced or supplemented by synthetic data. Moreover, data is only relevant in relation to specific applications. For instance, China’s data pool gives it an unsurpassed advantage in understanding Chinese consumer habits, but it may not confer wider advantage.⁴⁹ Russia’s battlefield testing of robotic systems on behalf of a ruthless dictator in Syria and the Chinese Communist Party’s creation of a massive AI-enabled surveillance state are signs of governments that fear their own people, do not trust their soldiers, and seek technical solutions to centralize power.

THE STATE OF AI IN GOVERNMENT

The U.S. government has issued strategic guidance that acknowledges the centrality of AI for national security, starting with the White House’s 2019 Executive Order on Maintaining American Leadership in Artificial Intelligence.⁵⁰ Members of Congress have filed over 30 bills addressing AI over the last five years, and organized a Congressional Artificial Intelligence Caucus. Senior leaders are beginning to appreciate AI’s impact on their organizations. The military services are developing operating concepts that account for AI and automation. Hundreds of promising AI-related projects are underway in all corners of DoD.⁵¹ A new entity, the Joint Artificial Intelligence Center, is entering its second year. Project Maven offers an example of a narrowly focused use of AI to detect, classify, and track objects on video streams so human analysts do not have to stare at screens for hours on end. The Intelligence Community has launched its own AI initiative.⁵² The Department of Energy has established an AI office to coordinate its efforts and to ensure that AI researchers have access to government data models and high-performance computing resources.⁵³ These are all positive signs of government organizations waking to AI’s potential. Most government officials the Commission has spoken to, however, recognize that their organizations remain far from where they need to be. The Commission believes the U.S. government still confronts enormous work before it can transition AI from a promising technological novelty into a mature technology integrated into core national security missions.

To a large degree, AI remains a hard problem for the U.S. government because AI does not fit the traditional paradigm of a technological development driven by national security needs and federal dollars. The opposite is occurring, with the DoD adopting a commercial technology for military use. Most AI applications currently are in the commercial sector, while the national security uses mostly reside in the realm of theory and possibility. Federal research support is critical to early stage AI research, but the government has limited control over how AI technologies are developed, shared, and used. Today's AI leaders and biggest funders can be found in universities, startups, and big tech firms.

The reversal of the Cold War paradigm where government R&D leadership spun out into the commercial sector poses serious obstacles to adopting AI for national security. In perpetual catch-up mode, the government is trying to integrate AI into existing infrastructure and technology, which is sometimes decades old. The government depends on the commercial sector, while the AI industry, far from depending on government business, often sees government regulations and bureaucracies as hindrances to their business models and therefore an unworthy pursuit. The government is trapped using an acquisition and testing and evaluation system designed for a different era that is ill-suited for AI's iterative, software-based attributes. Despite pockets of excellence, the government lacks wide expertise to envision the promise and implications of AI, translate vision into action, and develop the operating concepts for using AI. Finally, many of the gains from AI-enabled systems can only be realized through transformation of organizational structures and business processes; the inherent rigidity of government in this respect poses a major obstacle.

The challenge the government now faces is how to marry bottom-up and dispersed innovation with top-down vision. A vast space exists between commercial AI developments and field experimentation, on the one hand, and the broad AI guidance emerging from the top, on the other. The history of military innovation shows that organizational reform and overcoming adoption and deployment challenges will be as important as solving technical problems. Given the pace of AI developments and the actions of our competitors, the government must move faster.

In sum, the Commission's preliminary judgment is that the United States is not translating broad national AI strengths and AI strategy statements into specific national security advantages. Key departments and agencies have not yet fully embraced high-level strategy pronouncements and therefore critical national security missions have not incorporated AI. Given the general support for AI initiatives within government, one purpose of this interim assessment is to understand and explain the multiple reasons why that is so. We see the path to successful maximization of AI for national security purposes running

through five lines of effort. These are outlined below along with some preliminary initial judgments emphasizing specific areas that need more attention or may be ripe for action.

IV. *Lines of Effort for the U.S. Government*

1. Invest in AI Research and Development

The foundation of America’s global competitiveness in AI is dependent upon achieving technological breakthroughs in federal, academic, and commercial research and development.

Technology R&D in the United States has long been driven by a “triangular alliance” among government agencies, universities, and private companies. Created in the early days of the Cold War, that alliance propelled the country to global technological leadership and economic prosperity.⁵⁴ Defense research programs led to technologies, like stealth and the global positioning system (GPS), that were a generation ahead of what was commercially available or what our competitors could field. Many of these innovations—most notably the Internet and advanced microchips—fueled the rise of the tech sector we know today, giving U.S. companies an important lead over global competitors.

Since the Cold War, the triangular model has changed in important ways, including for AI. Much more R&D is happening within technology companies than ever before.⁵⁵ Federal R&D funding has continued a decades-long decline.⁵⁶ Academic-corporate collaborations in AI research are increasing.⁵⁷ While the commercial sector now plays a significant role in AI research, its investments are necessary but insufficient for sustaining U.S. advantages. The government retains a critical role, especially in supporting basic scientific research as well as research that is directly relevant to national security.⁵⁸ Like the transformational technologies that came before it, AI will reach its fullest potential when supported by government investments.

Despite the transformative potential of AI, the U.S. government has not yet responded with the resources necessary to meet current research needs and set conditions for future innovation. America’s AI leadership may be at risk sooner than is currently appreciated. Plainly stated, a loss of national leadership in AI technology development will mean that the U.S. military and intelligence agencies will have to acquire and use inferior systems—or buy better ones from China or elsewhere. We offer some initial judgments on the urgent need for additional federal R&D funding, new mechanisms for channeling those resources, and important focus areas for national security.

INITIAL CONSENSUS JUDGMENTS:

Federal R&D funding for AI has not kept pace with the revolutionary potential it holds or with aggressive investments by competitors. Investments that are multiple times greater than current levels are needed.

Requested FY 2020 federal funding for core AI research outside of the defense sector grew by less than 2 percent from estimated FY 2019 levels.⁵⁹ Over the past five years, federal R&D funding for computer science (which houses AI) increased by 12.7 percent,⁶⁰ barely sustaining a field in which tenure track positions grew by 118 percent over the same period.⁶¹

The U.S. government knows how to infuse resources into audacious technology projects, as it did for the Apollo space program or the Human Genome Project. While the Chinese government has made ambitious public commitments to technology megaprojects,⁶² the United States has returned to pre-Sputnik levels of federal R&D funding as a percentage of GDP.⁶³ Indeed, the trend is going in the wrong direction, with a proposed 5 percent cut to R&D funding (and 10 percent in basic research) in the FY 2020 budget.⁶⁴ The United States now trails nine nations on the measure of total R&D expenditure as a percentage of GDP.⁶⁵

Increased federal R&D investments could spur the development of:

- core AI technologies, such as unsupervised or self-supervised ML, AI systems with greater common sense, and AI inspired by neuroscience;
- AI systems that are safer, more robust, and resistant to attack;
- AI techniques to accelerate progress in important science and engineering problems, such as slashing the time needed to discover advanced materials;
- cloud infrastructure, labeled training data and other resources for AI researchers;
- next generation hardware, dedicated chips, and novel computing paradigms needed to fuel AI; and
- the workforce needed to develop and use AI effectively.

Limited availability of federal funding contributes to an accelerating brain drain from academia to industry.⁶⁶ This trend damages our ability to train the next generation and influences the direction of research toward more commercially-applied problems. The government must help redirect this trend soon.

The overall FY2020 budget request for non-defense federal AI R&D is \$973 million.⁶⁷ The National Science Foundation (NSF) is critical to the research system, as it manages 85 percent of all federal funding for academic computer science research in the United States. We have found that the NSF’s budget for basic AI research would need to double simply to cover only the most highly qualified proposals it receives through its rigorous peer-review process.⁶⁸

For maximal effect, federal investments should be phased across the near, medium, and long term. On the defense side, DARPA’s AI Next campaign is an example of a large-scale initiative that points in the right direction. Launched in 2018, it will provide \$2 billion in federal funding over five years. Grants are focused on defense-oriented problems, and what the agency calls “Wave 3” AI technology, in which machines do context-based reasoning and partner with humans.⁶⁹ Overall, about a third of DARPA’s current programs involve AI in some way.⁷⁰

Untapped opportunities exist to build a nationwide AI R&D infrastructure and encourage regional innovation “clusters.” Such AI districts for defense would benefit both national security and economic competitiveness.

We are considering a range of organizational models that could accelerate R&D nationwide. For example, the NSF has launched an effort to fund a series of AI R&D institutes, funded at \$4 million per year for five years (with an option to extend for another five). NSF expects to launch six institutes in 2020.⁷¹ In addition, we are examining other ideas, including establishing an entity within the NSF analogous to the National Cancer Institute, a structure resembling the National Institutes of Health to coordinate research and set standards, or an interagency AI effort akin to the National Nanotechnology Initiative.⁷² One proposal from a computing community consortium has called for a series of national AI research centers and labs, funded at \$100 million per year for at least ten years.⁷³

Economic trends suggest that people and firms will be drawn to geographic “clusters,” such as Silicon Valley.⁷⁴ With federal installations spread widely, the government has an opportunity to apply a broad regional lens to partnerships with academia and industry. Canada, for example, has taken a nationwide approach by spreading AI research centers across Edmonton, Toronto, and Montreal.⁷⁵ The new NavalX Tech Bridges program,⁷⁶ and the Navy Surface Warfare Center in Crane, Indiana,⁷⁷ are promising efforts to build such tech hubs.⁷⁸

The U.S. government should implement more flexible funding mechanisms to support AI research. Business as usual is insufficient.

The traditional model of short-term, project-based grants may be insufficient to foster transformational advances. We are exploring a number of alternative funding vehicles. For instance, mid-career faculty awards for AI researchers could encourage professors to remain in academia, rather than jump to industry, at a typically productive point in their careers. Subsidies to universities for AI degree development at the undergraduate and graduate levels, and for certifications in AI and advanced computing, could expand the talent pool. Expanded fellowships for graduate and postgraduate researchers would help develop more future professors.

Another promising idea involves more investments in individual lab leaders as is done, for example, at the Howard Hughes Medical Institute⁷⁹ and through DoD's Vannevar Bush Faculty Fellowships.⁸⁰ This approach prioritizes human capital, and gives top scientists the flexibility to adjust the direction of their research midstream, as they find new and promising avenues of inquiry.

The U.S. government can partner with the commercial sector to help AI researchers overcome the substantial technical and financial barriers that can hinder AI research.⁸¹ The government must continue to make its rich data sets available to researchers, foster public-private partnerships to provide cloud-based computing resources,⁸² and facilitate industry data repositories that can be accessed by verified researchers.⁸³

The U.S. government must identify, prioritize, coordinate and urgently implement national security-focused AI R&D investments.

The White House's 2019 Executive Order on Maintaining American Leadership in AI provides a sound policy framework, and recognizes the importance of AI R&D to create "capabilities that contribute to our economic and national security."⁸⁴ The National AI R&D Strategic Plan, prepared by the Office of Science and Technology Policy in 2016 and updated in 2019, has outlined a sensible list of goals.⁸⁵ The Networking and Information Technology Research and Development program has adopted a new framework to measure AI investments across civilian agencies.⁸⁶ We welcome these important steps, but they do not provide additional resources, and they stop short of spelling out how the government should focus its investments in areas that are most important for national security. A National Security Presidential Memorandum, "Protecting the United States Advantage in Artificial Intelligence and Related Critical Technologies," was signed in February, and the Commission looks forward to continuing its engagement with this effort.⁸⁷ Government investments in AI research should put a

premium on issues that are especially relevant to national security missions that the commercial sector may not have incentive to prioritize. These priorities should have high-level executive branch coordination.

For example, in defense and intelligence contexts, labeled data may be in short supply, requiring algorithms that can learn from limited or synthetic data. There is also a greater need for tactical computing, for operators who are deployed at a distance from centralized resources, in a contested environment with only intermittent communications links. Moreover, AI systems need to withstand a siege of adversarial attacks.⁸⁸ Remarkably, a very small percentage of current AI research goes toward defending AI systems against adversarial efforts.⁸⁹ To achieve true human-machine teaming in high-stakes situations, the operator must trust that the AI system functions properly. The degree of robustness of the AI systems that DoD and the IC need exceeds what is commonly available on the commercial market. To ensure that robustness, fundamental research into the science of validating AI technologies is critical. Finally, as deepfakes become more difficult to detect, research into digital forensics will become even more important.⁹⁰ The Commission will continue to develop these and other ideas for national security AI research priorities.⁹¹

Bureaucratic and resource constraints are hindering government-affiliated labs and research centers from reaching their full potential in AI R&D.

DoD's federally funded research and development centers (FFRDCs), where a lot of the most critical, mission-focused AI R&D is happening, are limited by legislative caps on funding and staffing.⁹² The Government Accountability Office found that the current ceiling "significantly constrains" DoD's use of these research centers, and that demand for their services is "significantly greater" than what the legislation allows. Department officials reported that "FFRDC related work must be deferred to later years when these limits are reached, since there are no other legally compliant alternatives capable of fulfilling these requirements."⁹³

Moreover, we have found that red tape in the DoD-owned lab network slows its ability to innovate. Layers of management and long approval processes lead researchers to choose older hardware and software for their work, because these can be obtained more quickly than the best products available. Such issues are creating risks that DoD labs will fall behind the curve of current AI research and development.

2. Apply AI to National Security Missions

Technological change is creating first-order national security challenges for the United States. Strategic competitors have invested in advanced capabilities to erode American military advantages, exploit vulnerabilities, and undermine conventional deterrence.⁹⁴ Decisions about how to apply AI for national security purposes will help determine whether the United States reverses the trends and meets its traditional national security objectives: defending the homeland, deterring war, protecting allies, and winning on the battlefield.

U.S. military strategy has long relied on technological innovation to achieve strategic objectives and maintain U.S. advantage. The past two decisive military-technological revolutions, known as the “first offset” and “second offset,” were enabled by specific technological breakthroughs that solved core defense problems. The first offset began early in the Cold War and solved the problem of how to deter a massive conventional Soviet threat to Western Europe by miniaturizing nuclear components to build a new generation of battlefield atomic weapons. Once the Soviet Union achieved nuclear parity, the second offset of the 1970s and ’80s sought to renew U.S. advantage and restore a credible deterrent by exploiting innovations in information technologies and digital microprocessors. These efforts led to new conventional guided weapons, stealth, GPS, and wide-area ground surveillance—technological innovations that enabled the United States and its allies to “look deep and shoot deep.” Each offset drove the development of new battlefield capabilities able to resolve a core defense dilemma.⁹⁵

Today, strategic competitors have caught up with the United States technologically, and threaten U.S. military-technical superiority. Efforts to address this development began with the 2014 Defense Innovation Initiative, which paved the way for a third offset strategy, the precepts of which were reflected in the 2018 National Defense Strategy.⁹⁶ The Commission believes AI is key to the next technological leap which, if leveraged appropriately, will equip the United States to extend its advantages and preserve a credible deterrent in East Asia and Eastern Europe. AI-enabled systems could allow U.S. forces to understand the battlespace more clearly and rapidly; make better informed and faster command decisions; use autonomous systems to mount operations even when communication links are under attack; and develop capabilities to better defend against adversary AI systems. Intelligence agencies will be able to integrate massive amounts of data and better identify threats and discern patterns, which will provide military commanders and policy makers with more timely and sophisticated analysis.⁹⁷

AI can enable our national security agencies to understand, operate, and execute their missions faster. However, the DoD and the IC are still a long way from realizing AI’s

potential benefits. Efforts to integrate AI face obstacles throughout the adoption pipeline. The history of military innovation suggests that overcoming these challenges is essential for effectively using new technology.⁹⁸ As with difficulties incorporating mechanized armor, steam-powered ships, aviation, and other technological advances into military operations, government can be its own worst enemy as it tries to translate AI strategy on paper into practice and technological breakthroughs in the lab into results in the field.

The challenge is compounded because the government has not yet effectively incorporated commercial developments or found a way to be a fast adopter of the latest technologies. To remain competitive, the U.S. government must accelerate efforts to apply AI and rethink military doctrine, strategy, organization, budgeting, acquisition, talent management, tactics, training, and infrastructure.

INITIAL CONSENSUS JUDGMENTS:

AI can help the United States execute core national security missions, if we let it.

With better AI applications, DoD can take advantage of autonomous and intelligent systems that can help our forces become more effective; the Intelligence Community can more effectively process and analyze vast amounts of data; and agencies can find efficiencies in business operations, so those resources can be reallocated to the highest priority missions. Across these contexts, AI applications enable greater autonomy, automation, speed, endurance, scaling, information superiority, and decision-making.⁹⁹

The potential value for national security missions is significant and wide ranging. For example, AI can process information and react at superhuman speed, providing an advantage in missions where speed is critical, such as cybersecurity or missile defense. In electronic warfare, cognitive systems could autonomously detect and respond to signals jamming.¹⁰⁰ AI-enabled autonomous systems can also operate with superhuman endurance, providing, for example, around-the-clock overhead reconnaissance. In anti-submarine warfare, an unmanned vessel could navigate the open sea and hunt adversary submarines for months at a time.¹⁰¹ And AI can help scan vast quantities of data to provide options to decision-makers, about, for example, prioritizing maintenance needs or selecting which forces and equipment to send into battle.

Implementation of the government's security strategies for AI is threatened by bureaucratic impediments and inertia. Defense and intelligence agencies must urgently accelerate their efforts.

On paper, the government clearly acknowledges the importance of AI for national security. The National Security Strategy, the National Defense Strategy, DoD's AI and Digital Modernization strategies, and the Intelligence Community's AIM Initiative all recognize AI as a transformative technology.¹⁰² Secretary of Defense Mark Esper identified AI as a top modernization priority.¹⁰³ However, it is not clear that these top-level beliefs and strategic priorities have been fully embraced by departments and agencies. There must be broad organizational understanding of how AI can address core national security challenges and what is needed to achieve an AI advantage. Without clear communication linking vision to organizational change, adoption will stall and AI could be consigned to a series of niche applications, or dismissed by skeptics as the next tech fad to be waited out.

Pockets of successful bottom-up innovation exist across DoD and the IC. These isolated programs cannot translate into strategic change without top-down leadership to overcome organizational barriers.

Early adopters within DoD and the IC are leading the way with flexible programs that can help define requirements and demonstrate value. AI application efforts remain largely decentralized. A recent estimate suggested there are over 600 active AI projects across DoD.¹⁰⁴ The military services, intelligence agencies, government labs, and other components carry out programs according to their own policies, processes, and budgets. Many projects have been carried out through innovation hubs within the services that take advantage of flexible acquisition options or through Special Operations units with authorities unique to that community.¹⁰⁵ Each project is unique in how it was established, fielded, and managed. Though the projects are promising, DoD is struggling to shift bottom-up experiments into established programs of record. Individual programs are not creating a critical mass for organizational change.

Building on Project Maven's work as an AI pathfinder, DoD created the Joint AI Center (JAIC) to help bridge bottom-up programs and top-down leadership. Its operating model affirms that most tactical AI innovation should still be developed at the edge. However, centralized direction and a common foundation can facilitate decentralized development and experimentation. Established in June 2018 under DoD's Chief Information Officer, the JAIC's objective is to accelerate delivery of AI-enabled capabilities, scale the DoD-wide impact of AI, and synchronize DoD's AI activities. The center is working with the

services and other components to identify new mission initiatives, rapidly deliver prototypes that show value, and build momentum for AI across DoD. Initial projects include applying AI to logistics, disaster relief, cyber operations, maneuver and fires, automated business processes, and service member health care.¹⁰⁶ The JAIC also provides a top-down coordinating role for DoD AI initiatives that have budgets exceeding \$15 million.¹⁰⁷

The IC's AI initiatives have benefitted from a concerted coordination effort and substantial senior leadership support, combined with investments in required infrastructure. Applications to date have focused on collecting, processing, and analyzing data to assist both analysts and operators. The 2019 AIM Strategy provides the framework for incorporating AI, process automation, and IC officer augmentation technologies across the community.¹⁰⁸ The strategy lays out a clear path forward with time-based objectives focused on furthering the IC's digital infrastructure and data systems, adopting commercial products, and investing in "sensemaking" research and technologies.¹⁰⁹

Like past technological changes, integrating AI will require top-down leadership and effective coordination to overcome cultural, policy, and process barriers to adoption. Direct attention from senior leadership can drive organizational focus and empower coordinating entities to avoid duplication of effort, share lessons learned, and scale AI programs.

| *AI adoption and deployment requires a different approach to acquisition.*

The government must develop some applications in house, while rapidly adopting and modifying commercial technologies. The current acquisition system was designed for big hardware programs with long, linear development timelines. AI and other software require more rapid, flexible, and iterative approaches.¹¹⁰ Conversations across the JAIC, the services, labs, and the IC indicate that delays in fielding AI are almost always caused by policy and process constraints within the acquisition system, not technical barriers. For example, an application might take two weeks to develop, but require over six months to get approvals to share the necessary data, another six months to a year to receive different approvals to run the application on an existing platform or network, and yet another several months to receive software updates. This ponderous process reflects a peacetime mentality and risk-averse culture that will hinder AI acquisition. It is part of a bigger software acquisition deficiency that inhibits DoD and the IC from fielding AI technologies.¹¹¹

In addition, DoD is struggling to access the best AI technology on the commercial market. Many leading commercial AI firms are small, far from Washington, D.C., and do

not consider the government market a priority. Traditional acquisition approaches have also made DoD an unattractive business partner for many top commercial AI firms, especially small and medium-sized businesses. We have identified a series of four common challenges DoD faces when working with tech companies ([Appendix 2](#)). Further research by the Commission will focus on recommendations to overcome these barriers and expand the number of top commercial AI firms engaging with the Pentagon and intelligence agencies.

DoD has developed some innovative approaches, including promising Air Force partnerships with MIT and Army partnerships with Carnegie Mellon University.¹¹² The services also utilize Rapid Capabilities Offices, and the Defense Innovation Unit's budget has grown.¹¹³ The problem is that such efforts have not been scaled up, and DoD continues to over-rely on niche organizations and institutional workarounds. Until the government overcomes these challenges, it will miss out on timely access to cutting-edge commercial breakthroughs and top AI talent.

Rapidly fielding AI is an operational necessity. To get there requires investment in resilient, robust, reliable, and secure AI systems.

There is a tension between fielding applications as quickly as possible and ensuring they perform reliably and safely. In finding the balance, we must not allow technical hurdles to serve as excuses for inaction. AI applications require iterative testing, evaluation, verification, and validation (TEVV) that incorporates user feedback and helps the system learn and improve. As a result, there can be a benefit to fielding applications early in their development cycle. TEVV requirements are even more important and difficult to meet for national security applications that must operate in complex adversarial environments and with limited training data. To minimize risks and accelerate fielding, DoD and the IC must invest in understanding and addressing technical vulnerabilities throughout the development of AI systems, including from organic operational errors and adversarial attacks. These prerequisites for fielding trustworthy and resilient AI applications can also serve as a basis for developing U.S. counter-AI and counter-autonomous capabilities against adversary systems.¹¹⁴

AI is only as good as the infrastructure behind it. Within DoD in particular this infrastructure is severely underdeveloped.

Even the most advanced AI algorithms will fail if they are not supported by sufficient computing power, data, storage, and communication networks. Significant investments in this foundational infrastructure are necessary across DoD and the IC, especially for data management and tactical edge applications. Modernizing DoD's IT infrastructure

will require significant investments in, among other things, the cloud and computing platforms necessary for data storage, compute resources, network communications, and algorithm development.

With respect to data, the government is well positioned to collect useful information from its worldwide network of sensors. But much of that data is unlabeled, hidden in various silos across disparate networks, or inaccessible to the government because of contracting regulations. Even more data is simply expelled as “exhaust” because it is not deemed to be immediately relevant. DoD strategies recognize the importance of treating data as a “strategic asset;” now, the challenge is to build the needed infrastructure and adopt new routines that foster an ecosystem for data success.¹¹⁵

In some cases, the government suffers from a data deficit. For example, when intelligence agencies train an algorithm to find specific objects, there may be very few images available to use in the training. As a result, there is a need for synthetic data and for algorithms that can learn from limited data sets. In addition, national security agencies face unique challenges in managing data across security classification domains. Methods to use unclassified surrogate data, homomorphic encryption, or other means to train algorithms for use on classified programs should be prioritized. DoD and the IC will also need to invest in purpose-built edge processing and network architectures for AI applications in forward operating missions, including in harsh environments where communications may be disrupted.

The U.S. government is not adequately leveraging basic commercial AI to improve business practices and save taxpayer dollars. Departments and agencies must modernize to become more effective and cost-efficient.

Proven off-the-shelf commercial AI solutions can make back-end processes such as human resources, financial management, contracting, and logistics more efficient and cost-effective across large organizations. These applications enable automation of repetitive, high-volume tasks, saving time and money and freeing workers to concentrate on the human dimensions of their work or support the mission in different ways.

3. Train and Recruit AI Talent

In a strategic competition, advantage will go to the competitor that can best attract, train, and retain a world-class, AI-ready workforce. Currently, there is a severe shortage of AI knowledge in the DoD and other parts of government.

Our defense and intelligence agencies need access to more people with AI skills and expertise, both in-house and outside of government. Their knowledge is required to buy, build, and use AI tools effectively. Aside from small pockets, however, the government has been slow to recognize the importance of technical skills needed for an AI-ready workforce, and has struggled to attract and retain AI talent.

Because of this deficit, the government is struggling to implement AI solutions. Ineffective talent recruitment and management, combined with outdated data management practices and a lack of processing power, create obstacles to integrating AI systems. As a result, projects rack up higher costs and require longer timelines for innovation and deployment.¹¹⁶

An AI-literate federal workforce that understands the basics of computational thinking, when and how to purchase commercial AI tools, and how to develop and implement custom systems can catalyze government AI use. Without investments in expertise and broader familiarity with the technology, AI adoption will be inconsistent at best, and the promise of an AI-enabled future will be compromised.

As with any imperative to increase human capital, the government can train the people it has, recruit new talent, and partner with outside institutions. Rethinking the AI workforce requires that we consider everyone from sergeants in the field to postdocs in the lab.¹¹⁷

The Commission is looking into the state of the nation's AI talent pool—from which the military and national security agencies will draw their expertise—and the challenges of recruiting for public service in today's labor market. The Commission has repeatedly heard from industry and government that AI experts would be willing to serve in government if officials could create a more compelling sense of purpose and a technical environment within government that would maximize their talents. This inquiry will also lead us to examine relevant aspects of the U.S. education system and the role of international students and workers.

INITIAL CONSENSUS JUDGMENTS:

National security agencies need to rethink the requirements for an AI-ready workforce. That includes extending familiarity with a range of relevant AI technologies throughout organizations, infusing training on the ethical and responsible development and fielding of AI at every level, and spreading the use of modern software tools.

National security organizations must have AI workforces capable of performing six functions: 1) planning and executing an organization-wide strategy; 2) purchasing and maintaining software and hardware infrastructure; 3) managing and analyzing data; 4) developing software when necessary, for unique needs; 5) performing verification, validation, testing, and evaluation; and 6) deciding when and how to employ AI tools. Given these requirements, an AI-ready workforce must include a solid nucleus of AI technical experts and developers. But the bulk of the workforce will be people who enable, or are enabled by, the effective use of AI. This larger group needs to understand the fundamentals of AI policy, functionality, and application. Accordingly, for DoD and many intelligence agencies, familiarity with AI should be more widespread, from senior leaders to mid-level officials to technical staff.¹¹⁸ Developing AI-ready leaders is especially critical. Without more well-informed leaders who can go beyond talking points and reshape their organizations, the defense and intelligence communities will fail to compete in the AI era. We have drafted a typology of organizational roles and associated AI knowledge and skills ([Appendix 3](#)).¹¹⁹

We offer two additional observations. First, an understanding of ethical dimensions of AI should cut across all organizational roles. Ethical questions can arise at every stage of AI development and use, so broadening a familiarity with these risks and concerns throughout our military and civilian institutions is a priority. Second, the technical workforce needs the right tools, authorities, and office environment to be effective. They need access to sufficient computing, storage, and data; adequate authorities for state-of-the-art systems development; an efficient software purchasing process; and an office space organized for deployment of these tools.¹²⁰

DoD and the IC are failing to capitalize on existing technical talent because they do not have effective ways to identify AI-relevant skills already present in their workforce. They should systematically measure and incentivize the development of those skills.

Four near-term actions can help. First, the military reserve components should develop a tracking system for civilian employment and accompanying skills, to better understand what skills are already in the reserve forces.¹²¹ Second, DoD and the IC should reward employees for learning AI-relevant skills by paying bonuses for completing courses. The U.K. Royal Air Force has introduced a pilot program along these lines.¹²² Third, the Armed Services Vocational Aptitude Battery Test should include measurements of computational thinking required for AI software development proficiency. The Air Force is currently exploring integration of a computer language aptitude test to identify potential developers.¹²³ Finally, Services should treat coding like a foreign language—by allowing service members to test for proficiency, and rewarding proficiency with incentive pay. The Air Force is already developing such a program.¹²⁴

The U.S. government is not fully utilizing civilian hiring authorities to recruit AI talent. Agencies need to make better use of pipelines for people with STEM training.

The U.S. government has the vast majority of the authorities it needs to hire AI talent. Several authorities are especially relevant to improving the AI workforce. Agencies with a critical need can shorten the hiring process, using authorities specific to STEM fields, especially direct hiring authorities and excepted service.¹²⁵ But many organizations underutilize these authorities, often due to risk-averse human resources teams that play an outsized role in the vetting process and often rely on imprecise position descriptions rather than an understanding of the position's technical requirements. Agencies can also hire consultants and highly qualified experts outside of the competitive service system.¹²⁶ Several agencies use the Pathways Program and other internships as hiring pipelines.¹²⁷ The Presidential Management Fellows program is a central pathway into government for graduate students.¹²⁸ Finally, DoD and other agencies run several scholarship programs with service obligations, such as Science, Mathematics, and Research for Transformation (SMART) and Cyber Corps (Scholarship for Service).¹²⁹ Any one of these programs, if focused and expanded, can increase the AI talent in government.

Expanding AI-focused fellowships and exchange opportunities can give officials and service members access to cutting-edge technology, and bring talent from our top AI companies into federal service.

The military sends a small percentage of uniformed officers to training with industry programs, and DoD and other agencies also encourage civilian rotations in companies, typically for one or two years.¹³⁰ These could be increased and focused on AI-relevant fields. National security agencies also need to work through the challenges of bringing in talent from the private sector for short-term assignments. Even where programs exist, issues with security clearances, intellectual property, and non-disclosure agreements create significant challenges.¹³¹ Congress has considered legislation to facilitate public-private exchanges in cyber security, which could be used to cover AI and related fields.¹³² To bring in more AI talent from universities and think tanks, the Intergovernmental Personnel Act could also be utilized more effectively.¹³³

The military and national security agencies are struggling to compete for top AI talent. They need a better pitch, incentive structure, and better on-ramps for recent graduates.

There is a tight labor market for AI talent in the United States, with demand far exceeding supply. Private sector salaries are often several times higher than what the government can offer based on today's pay scales.¹³⁴ But the government needs to bring in a new generation—not only to utilize AI, but also to analyze how other countries are using it.¹³⁵

Making the case for public service as a mission has become increasingly difficult. The NSA's general counsel lamented that “millennials believ[e] that technology in the private sector now allows them to help change the world,” whereas “previously the idea of a mission had largely been the province of public service.”¹³⁶ Still, we believe the government can improve its recruitment pitch by demonstrating that there are compelling ways to contribute to society while solving uniquely challenging, important problems. The government can create new pathways into meaningful service—through, for example, a technology-focused fellowship program for recent college graduates.¹³⁷ Some have also called for a U.S. Digital Service Academy and a Reserve Officer Training Corps (ROTC) program focused on advanced technologies.¹³⁸

Moreover, STEM fields, including AI, have a well-documented diversity problem.¹³⁹ Without change, the government will not harness the full potential of the American people or the promise of the technology.¹⁴⁰

Two realities about the American AI talent pool have become clear to us:

- 1. Colleges and universities cannot meet the demand for undergraduate student interest in AI and computer science generally.*

Many computer science departments are struggling with huge increases in the number of students, without corresponding increases in faculty and resources to properly serve them. Universities urgently need to hire more faculty to cover the teaching load. A 2016 analysis of tenure-track computer science positions found that openings had jumped a staggering 71 percent in two years.¹⁴¹ In 2018, the Computing Research Association found that the number of computer science majors is increasing at ten times the rate of tenure-track faculty.¹⁴² But hiring is difficult because of the drain of AI researchers to industry and limitations in available research funding. These deficiencies are weakening the U.S. academic environment and degrading our ability to train the next generation in AI and related fields.¹⁴³ By contrast, Chinese universities are creating around 400 AI-related majors in 2019.¹⁴⁴ If unaddressed, the weakening U.S. academic environment in AI will have long-term, detrimental effects on the nation’s ability to harness AI.

The Commission will study the problems at the undergraduate and graduate level, along with other components of the U.S. education system—including post-secondary credentialing opportunities for non-college educated workers, the state of K-12 education in preparing students for higher education in AI-related fields, and the potential to rapidly expand online and other non-traditional AI graduate programs.¹⁴⁵

- 2. The American AI talent pool depends heavily on international students and workers. Our global competitiveness hinges on our ability to attract and retain top minds from around the world.*

Worldwide, the supply of AI talent is insufficient to meet a growing demand. Given this scarcity, talent is among the most valuable inputs into a nation’s AI ecosystem.¹⁴⁶ A recent study found that “the majority of workers in AI-related jobs and students in AI-related graduate programs are not originally from the United States.”¹⁴⁷ AI is a highly mobile field, especially among PhD-level experts. The 2019 Global AI Talent Report determined that about one third of AI researchers work for an employer based in a country other than the one where they received their PhD.¹⁴⁸ The United States remains the destination of choice for the largest number of internationally mobile students and researchers in AI,¹⁴⁹ and should ensure its lead by strengthening high-skilled immigration opportunities for the world’s most talented scientists and engineers.¹⁵⁰

4. Protect and Build Upon U.S. Technology Advantages

For decades, the United States has maintained an open economy and championed academic freedom, while also protecting its edge in defense and security-related technologies. It has preserved this balance through robust counter-intelligence, reviews of foreign investment, and export controls, among other techniques. Those tools remain important. But certain features of the current geopolitical and technology landscape are straining America's ability to institute a coherent and effective technology protection regime:

- The nature of AI technologies makes the protection of those technologies for national security very difficult. AI research has been largely decentralized and industry-driven; as a result, knowledge is more diffuse and accessible than historical breakthrough technologies such as nuclear or stealth.
- Open access to AI research is a strong norm in computer science. Even if restrictions were placed on AI products or services, much of the underlying code is publicly available.¹⁵¹
- The United States and China have close linkages in the field of AI, including constant exchanges of people, research, and funding. Chinese AI researchers train at U.S. universities. American cities host Chinese AI research centers, and major U.S. companies have research ventures in China. Chinese venture capitalists have invested in American AI start-ups, and vice versa.
- At the same time, China takes advantage of the openness of U.S. society in numerous ways—some legal, some not—to transfer AI know-how.¹⁵² U.S. intelligence agencies confirm that the “targeting of national security information and proprietary technology from U.S. companies and research institutions will remain a sophisticated and persistent threat.”¹⁵³
- America's research universities thrive by welcoming top minds from around the globe. At the same time, universities and other research institutes are vulnerable to foreign exploitation and other forms of influence by strategic competitors, notably China.¹⁵⁴

Taken together, it is difficult to resolve these issues in ways that balance security concerns with the principles of an open society and the core American traditions of free enterprise and free inquiry. Our initial research has led us to four broad judgments.

INITIAL CONSENSUS JUDGMENTS:

The U.S. government should continue to use export controls—including multilateral controls—to protect specific U.S. and allied AI hardware advantages, in particular those in semiconductor manufacturing equipment.

AI applications rely on hardware, and currently that hardware is almost exclusively powered by semiconductors.¹⁵⁵ Generally, countries with greater access to high-end computer chips will have an inherent advantage in their ability to deploy high-performing AI algorithms. The demand for semiconductors to enable AI applications is expected to grow dramatically in the coming years. U.S.-headquartered firms account for nearly half of all global semiconductor production.¹⁵⁶

As AI becomes more widespread and advanced, demand for more sophisticated and specialized chipsets to run algorithms will increase. This, in turn, will also increase demand for semiconductor manufacturing equipment (SME).¹⁵⁷ Due to the high cost and deep expertise necessary to construct SME, especially the most complex SME, this technology is heavily concentrated. About 90 percent of the SME industry is located in the United States, Japan, and the Netherlands, giving that small group of allies a major advantage.¹⁵⁸

Controls to preserve U.S. and allied advantages in SME could ensure that U.S. and allied country firms retain a dominant position in the global semiconductor market, including in advanced hardware capabilities.¹⁵⁹ It would also ensure that the U.S. government maintains access to the most cutting-edge hardware for AI applications and can scale up production in the event of a crisis. These steps would help to secure the global supply chain, protecting the United States and its allies from competitors' attempts to disrupt U.S. application of AI. It is critical that any such controls are multilateral in nature, as unilateral controls would negatively impact U.S. businesses, push R&D outside of the United States, and not measurably impact adversaries' ability to procure equipment.

At the same time, officials developing any new controls or regulations should consider the unintended consequences and financial hardship that might be imposed on U.S. companies. To the extent that U.S. companies could lose access to important parts of their global supply chains and markets, U.S. economic competitiveness could be harmed. The U.S. government may need to pair SME export controls with increased R&D investment in semiconductor design, manufacturing, packing, and testing in order to preserve continued American leadership in next-generation hardware.

Traditional item-based export controls and narrowly-scoped foreign investment reviews are by themselves insufficient to sustain U.S. competitiveness in AI.

The Commerce Department is leading a process to identify potential controls on “emerging” and “foundational” technologies considered essential to national security—including on AI technologies like computer vision, speech recognition, and natural language processing.¹⁶⁰ However, the multi-use nature of AI will not fit neatly with the current item-based approach to export controls. Instead of the traditional item-based approach, the government should consider heavier scrutiny of the potential end use and end user of specific items, to prevent their use for malicious purposes. Such a process is more resource intensive, but likely necessary to ensure that controls are effective and do not place an undue burden on U.S. corporations.¹⁶¹ The Commission will continue to examine which kinds of AI systems or components could lend themselves to control measures under such a framework, while keeping in mind the difficulties in using such controls, especially for software.

The past decade has seen an explosion of Chinese investment in U.S. AI companies: from \$1.5 million in just one deal in 2010, to \$514.6 million across 27 deals in 2017.¹⁶² The recent Committee on Foreign Investment in the United States (CFIUS) reform legislation, known as the Foreign Investment Risk Review Modernization Act, made important changes to the investment review process.¹⁶³ Continued implementation is necessary, and CFIUS should consider establishing a permanent review structure for AI-related investments. The Treasury and State Departments should also continue working with allies and partners to develop their own investment screening programs to prevent adversaries from migrating malicious investment strategies from U.S. to allied markets.

The United States must continue leading in AI-related hardware, and ensure the government has trusted access to the latest technologies.

It is critical that DoD and the IC retain access to trusted semiconductors and SME. Some trends are concerning.¹⁶⁴ DoD’s existing access to trusted hardware trails several generations behind commercial state of the art. Two recently launched efforts, the Microelectronics Innovation for National Security program and DARPA’s Electronics Resurgence Initiative, are revitalizing the U.S. government’s approach to microelectronics.¹⁶⁵ The Commission looks forward to engaging with these programs.

China has established a National Integrated Circuit Investment Fund to improve its hardware industry and increase self-sufficiency, investing over \$100 billion in the next decade.¹⁶⁶ The U.S. government will need to pursue advanced chip design capabilities to stay ahead in the global semiconductor market.¹⁶⁷ Over the longer term, the government

will need to focus on next-generation and potentially disruptive hardware solutions—such as silicon photonics and quantum computing—as those come into broader use in future AI systems. The Commission will examine these trends and what measures the government should take to leverage them.

Law enforcement and academic leaders can and should find common ground on preserving an open research system while reducing security risks from foreign government-directed activity on American campuses.

We fully recognize the threats to the integrity of our university communities and the research process posed by state-directed efforts.¹⁶⁸ At the same time, we affirm our shared belief that U.S. openness is among our nation’s greatest strengths, as is our ability to attract the best and brightest from across the world to contribute to innovation and discovery. We also take seriously the growing concerns about imprecise law enforcement and counterintelligence investigations, which can create the impression that certain groups of people are under heightened suspicion on the basis of their ethnic or national origin. In May 2019, the Administration established the Joint Committee on the Research Environment to examine research security.¹⁶⁹

Although universities are not intelligence or law enforcement agencies, they need to be part of the solution. Recent initiatives by major universities to scrutinize foreign influence are positive steps.¹⁷⁰ But university leaders have told us they are looking for clearer “guide posts” from the federal government on foreign collaboration risks and malign efforts directed by foreign governments. For its part, the government should help universities with awareness and capacity—by providing more declassified information with actionable guidance, as well as additional resources to pursue the necessary due diligence. The most focused attention should be placed on preventing direct or indirect assistance to China’s military and intelligence apparatus.¹⁷¹

The Commission is examining a number of ideas, including the role of an interagency task force on academic espionage, as proposed in recent legislation;¹⁷² heightened scrutiny during the visa process for Chinese researchers with certain risk indicators, such as ties to the Chinese military; security classification options for federally-funded AI research programs;¹⁷³ options for foreign students to remain in the United States after completing their studies; and ways for universities and the IC to have a constructive dialogue about potential threats.¹⁷⁴

5. Marshal Global AI Cooperation

The United States and like-minded nations must assume AI leadership now. Existing cooperative efforts between the United States and traditional and non-traditional partners provide a glimpse of the potential for closer AI partnerships between militaries, intelligence services, diplomats, and scientific researchers.

In the military and intelligence realm, AI alliances and partnerships can provide a framework for cooperative planning, data sharing, procurement, and interoperability across operational environments. A number of U.S. allies are integrating AI technologies into their military and intelligence platforms. In R&D, partners have launched programs devoted to new AI investments and research agendas. And the private sectors of U.S. allies and partners feature innovative AI companies and deep venture capital markets that are already providing capabilities to DoD and the IC.¹⁷⁵ The U.S. government's challenge will be to overcome significant technical, legal, and organizational barriers that stand in the way of cooperating—even with close allies.¹⁷⁶

Perhaps most significantly, the United States and its partners are in a competition to shape AI norms and use worldwide. The state or group of states that achieves technical leadership will have unique opportunities to set standards, build guard rails, and generate global support for what is acceptable and what is not in AI's future. As AI becomes a focus of multilateral bodies like the United Nations and the Organization for Economic Cooperation and Development, the Commission is considering ways to build coalitions that can advance U.S. and allied interests and values, particularly with regard to AI norms and standards.¹⁷⁷

We are also considering possible avenues for AI-related diplomatic discussions with rivals such as China and Russia, in areas such as AI safety and AI's implications for strategic stability, where we may be able to find common interests, promote responsible research and innovation, and limit dangerous uses.

INITIAL CONSENSUS JUDGMENTS:

The United States must enhance its competitiveness in AI by establishing a network of partners dedicated to AI data sharing, R&D coordination, capacity building, and talent exchanges.

The United States should aim to establish a network of like-minded nations dedicated to collectively building AI expertise and capabilities. This could include more coordinated AI R&D spending, and cooperative arrangements in data sharing, hardware, export controls, and talent exchanges, as well as efforts to improve AI literacy and computer science education. Coordinated AI R&D spending could more efficiently allocate allied resources, and pooling data centers and computing resources could increase collective AI capacity. The U.S. government also needs new approaches to assisting partners with AI adoption using traditional security assistance programs and other capacity building initiatives. To enhance collective competitiveness, the United States and its partners need to lower the barriers to the movement of people and data among nations, especially in light of China's enormous data and human resources.¹⁷⁸

Allies and partners have told us they are interested in continuing to develop common standards for ethical AI, including in areas such as data sharing, safety, and certification systems for trust and transparency. However, divergent views on data privacy present significant hurdles, in particular with respect to the European Union's General Data Protection Regulation. The Commission will explore the implications for AI cooperation in greater depth.¹⁷⁹

Strong U.S.-led diplomacy will play a critical role. The government should organize itself as soon as possible to conduct a sustained, long-term diplomatic campaign to support America's AI agenda. The Commission is aware of plans for the State Department to establish a bureau focused on cyberspace security and emerging technologies, and believes it should be done without further delay.¹⁸⁰

AI presents significant challenges for military interoperability. If the United States and its allies do not coordinate early and often on AI-enabled capabilities, the effectiveness of our military coalitions will suffer.

Close allies have informed us that they are concerned about being able to operate effectively together as the United States fields greater numbers of autonomous systems. To lay a foundation for enduring compatibility in the AI era, the U.S. government will

need to utilize existing science and technology cooperation agreements, and draft new, AI-specific data-sharing frameworks.¹⁸¹ The government also needs to ensure its data management, communications infrastructure, and information sharing authorities allow defense and intelligence agencies to connect with the networks of their foreign partners. The United States and its allies will need to rethink how they classify and store data sets, developing ways to combine pools of information to fuel combined AI systems. Overcoming these hurdles can help U.S. and allied militaries incorporate AI into combined operational concepts and fielded systems.

The Five Eyes alliance is a good place to start, because the United States can leverage existing technical cooperation and information sharing agreements.¹⁸² The Five Eyes Technical Cooperation Program recently embarked on an AI Strategic Challenge, a three-year effort focused on AI applications for allied militaries.¹⁸³

A harder task is developing interoperable systems throughout the 29-member NATO alliance. Different militaries will integrate AI at different rates, and there is a risk of divergence. At the same time, there might be ways for countries with fewer resources to specialize in particular AI functions—related to cyber or predictive analytics, for example—that would help them add value to the alliance. As one study pointed out, it may be easier and more advantageous for some allies to identify an AI-related niche than to invest in expensive equipment like advanced fighter jets.¹⁸⁴ There are encouraging signs that NATO is starting to incorporate AI into its training regimen and its strategic analysis. An allied exercise in 2018 called Trident Juncture, for instance, included 20 AI-related experiments, including innovative methods to find and treat battlefield casualties.¹⁸⁵

U.S. diplomacy should be open to possible cooperation with China and Russia on promoting AI safety and managing AI's impact on strategic stability.

AI presents significant challenges for developing arms control agreements. At this stage, it is not clear to us what such agreements could control or how compliance could be verified. But the dangers posed by AI-enabled military systems for global stability warrant a serious consideration of possible means to limit risks through diplomacy.

One avenue for engagement may relate to nuclear command and control. We are considering, for example, whether the United States could seek agreements with China, Russia, and others on issues where there may be mutual strategic interest, such as prohibiting the use of AI to authorize the launch of nuclear weapons. Even if technical verification is unlikely, a normative understanding may still be worth exploring.

Another potential topic is AI safety. We are exploring ways to establish greater international consensus on building safe and trustworthy AI systems to minimize unintended consequences. Discussions of AI safety should begin by identifying areas in which AI poses unacceptable risks of escalation or loss of control.¹⁸⁶ The United States could start by facilitating track 1.5 or track 2 dialogues between U.S. and Chinese researchers to discuss AI safety problems, promote safety research, and sketch a shared agenda on robust AI systems.¹⁸⁷

The United States should lead in establishing a positive agenda for cooperation with all nations on AI advances that promise to benefit humanity.

The Commission is considering the role of scientific diplomacy in promoting global AI collaboration in areas of AI application that can alleviate human suffering and provide common goods. Promoting advances in applications for cancer treatment or disaster relief, for example, could be of great interest to all countries and would affirm that the United States is leading a positive AI agenda in service of humanity.

V. *Considerations on Ethical and Trustworthy AI*

The Commission views the ethical and responsible development and deployment of AI technologies as a priority that cuts across all of the Commission's lines of effort. Ethical considerations are an important facet of R&D, application, training, protection, and cooperation in AI. When referring to ethical and trustworthy AI in the context of national security, we identify three essential components: 1) the ethical design and development of trustworthy AI systems; 2) the ethical use of these systems; and 3) the preservation of applicable rights and liberties when using AI. Developing trustworthy AI systems is essential for operational integrity and adoption. It is closely connected to, and depends on, reliability, robustness, auditability, explainability, and fairness. From the earliest phase, systems should be designed with ethics in mind.¹⁸⁸ Ethical use concerns the circumstances in which it is appropriate to deploy AI. Finally, throughout their life cycles, ethical AI systems for national security will need to preserve individual rights and liberties as protected by law. In international contexts, this includes America's commitments to international humanitarian law and human rights.¹⁸⁹ We believe it is on this basis that American AI systems and those of our democratic allies will be distinguished from those of authoritarian regimes.¹⁹⁰

There are ongoing initiatives within DoD and the IC to develop principles of AI ethics.¹⁹¹ The Commission commends the Defense Innovation Board on its development of AI principles and urges other agencies to engage in a similar process. Interagency cooperation may help further the implementation of such principles. For example, DoD and the IC are engaged in a combined effort to establish and integrate best practices for machine learning that will be complementary across organizations. Principles are an essential first step, and should be followed by processes to enact them across each organization. Each agency's design and deployment of AI, as with other technologies, must align with America's democratic values and institutional values.¹⁹²

There is broad consensus that AI systems should be trustworthy, explainable (or at a minimum auditable), and free of unwanted bias.¹⁹³ On these issues, the Commission will seek to provide a clear way forward on how best to operationalize these concepts. For example, recommendations might include identifying gaps in current processes, providing guidance on needed policy and technical standards, and identifying areas where future R&D and workforce training is necessary.

There is, however, disagreement among diverse groups on other issues, such as the ethics of using AI in armed conflict (including autonomous weapons systems), and how to use AI systems domestically in ways that preserve civil rights and liberties.¹⁹⁴ On issues that remain fraught, the Commission will seek to provide a clear point of view, informed by the perspectives and concerns of diverse stakeholders.

Rights and Liberties: AI-enabled tools, like facial recognition technology, may have valuable national security applications, but can also be used for racial profiling, violations of privacy, and targeting of vulnerable populations.¹⁹⁵ Federal law enforcement agencies conducting national security investigations in the United States should only use AI in ways that are consistent with constitutionally-rooted principles of individual privacy, equal protection, non-discrimination, and due process. The United States has a robust legal tradition that seeks to balance free enterprise, individual liberties, and public safety. The task is to harmonize the attributes (and current shortcomings) of AI technology with existing legal and ethical frameworks in each of these domains to ensure that AI is responsibly used.¹⁹⁶

International Human Rights: In the international context, the Commission is deeply troubled by reports that China's AI-powered surveillance is aiding the state in persecuting Uighurs and other religious minority groups.¹⁹⁷ Moreover, we are concerned that American institutions have ties to Chinese companies building these systems.¹⁹⁸ With the export of surveillance technology to repressive regimes on the rise, the United States should continue taking steps to prevent U.S. entities from unknowingly abetting abuses through robust export controls, disclosure requirements, and sanctions when appropriate.¹⁹⁹ While governance and a state's human rights record are the most important factors in anticipating how a government will utilize AI surveillance systems, the Commission is aware that over half of advanced liberal democracies are users of AI surveillance, necessitating safeguards for responsible use writ large.²⁰⁰

The Practical Benefits of Ethical AI: It is also critical for our national security institutions to uphold the highest ethical standards for strategic and pragmatic reasons. A defense apparatus that more clearly articulates its strong ethical commitments will be better able to recruit talent and partner with the private sector. Ethical employers are more likely to retain talent and cultivate higher morale among employees. Common standards for ethical AI will also enable the United States to partner more effectively with international allies who share our values. Most importantly, the ethical and responsible handling of AI generates trust among citizens.

VI. “Associated Technologies”

Our mandate from Congress is to study AI and “associated technologies” as they relate to national security. As the Commission continues its research, we will devote more attention to AI’s position within a constellation of emerging technologies that both enable and build upon one another.²⁰¹ For the moment, we offer some initial impressions that will shape our future investigation.

The sustainment of a technological revolution powered by data and new machine learning techniques is contingent upon the co-evolution of related hardware and computing technology. Developments in AI amplify and reinforce other technologies (and vice versa), underscoring the importance of supporting progress across the board, rather than redirecting resources to AI away from other areas of computer science. Given the interdependencies, we intend to focus on the related technologies that we expect to have the greatest impact on AI and those security-related fields that are most likely to be transformed by AI. The Commission will consider associated technologies that fall into two categories: a) research that is adjacent and complementary to AI; and b) domain-specific applications of AI. The Commission will examine these related areas insofar as they help to elucidate the full picture of AI’s potential in national security.

Adjacent Emerging Technology: Data, Networks, and Compute

Reliable, secure, high-speed connectivity is critical to powering AI systems. Developments in network infrastructure, including fifth-generation cellular networks (5G) have the potential to dramatically alter the security environment as it evolves alongside AI. By improving speed, volume, and latency, 5G networks will significantly enhance information-sharing capabilities in both defense and commercial contexts.²⁰² 5G infrastructure will support proliferation of AI-powered technology, enabling remote device operation, autonomous vehicles, smart cities, and homes connected by the Internet of Things. With increasing capability to process machine learning algorithms locally on-device, 5G is poised to serve as the unifying network layer that connects AI-enabled systems. In turn, data generated on 5G networks can be leveraged to improve AI systems, giving the data-owners a significant advantage that rapidly compounds. Underlying both 5G and AI are highly customized and complex microelectronics, underscoring the need for the United States to remain competitive in hardware development while considering the implied trade-offs and potential downsides.²⁰³ While the numerous economic and geopolitical challenges raised by the introduction of 5G fall outside the scope of this

Commission, we will consider network capabilities and limitations insofar as they dictate AI competitiveness. For example, we will consider how battlefield and national security infrastructures might be modernized to accommodate the volume of data required to support AI systems.

Progress in AI has also gone hand-in-hand with progress in computing. Developments in supercomputing and novel computing paradigms may have significant implications for the development of AI. For example, the Department of Energy is currently building next-generation exascale supercomputers that will apply high performance computing and AI for scientific research.²⁰⁴ With exponential increases in the volume of data and complexity of neural networks, quantum computing is one example of a new approach to computing that may unlock new processing capabilities.²⁰⁵ Though quantum computers have not yet reached viability at scale, breakthroughs in quantum represent a step-change advantage over the computational capacity of today's best supercomputers.²⁰⁶ Quantum computing may also arm its owners with the ability to crack current encryption, jeopardizing secured communications and transactions.²⁰⁷

The federal government has taken several promising steps toward supporting progress in quantum,²⁰⁸ which is increasingly urgent given China's significant investments in this field.²⁰⁹ The Commission notes, however, that not all developments in quantum are directly related to AI, and that the field of quantum machine learning is still in its infancy. To date, few ideas have been proposed regarding methods to harness (even theoretically) quantum computing to speed up machine learning and inference. Given quantum's application to highly specific problems, it is not likely to serve as the next general purpose solution for high-performance computing. Though we anticipate and encourage meaningful research in this area, the Commission's attention will be focused on developments in an array of computing technologies that will support a robust AI ecosystem, rather than on specific fields of research.

Application-Specific Domains

AI promises to enable a range of domain-specific applications with security consequences. For example, AI methods have long been employed in biomedicine. Further developments combining AI technologies with life sciences is likely to usher in a new chapter in biotechnology. AI methods have great potential to dramatically enhance our understanding of complex biological systems, as well as accelerate developments in gene-editing, genomics, and synthetic biology.²¹⁰ It also presents the risk of weaponization, as China is reportedly exploring.²¹¹ AI will also have important implications for security-critical fields such as agriculture, materials manufacturing, financial markets, and transportation. The Commission will consider the risks associated with applied-AI across

sectors and research domains, focusing on those with the most direct implications for national security.

Many of the risks and challenges of AI resemble those presented by associated technologies, including: the complexities of dual-use; the risks of malicious repurposing; the challenges of restricting tech-transfer of open-source research; the dominance of the private sector; and the relevance of first-mover advantage. In light of these commonalities, the Commission's approach to AI can be used to inform strategies for harnessing the potential of other emerging technologies. The reforms, structures, processes, or funding mechanisms that this Commission might recommend in its final report could be repurposed or expanded to include related technologies.

Appendix 1: What Is Artificial Intelligence?

The term “artificial intelligence” covers a broad range of computer system abilities to perform tasks that otherwise would require human intelligence or other forms of intelligence observed in nature.²¹² There is no consensus on a singular definition of AI; rather, there are many useful definitions that serve various communities and purposes. The Association for the Advancement of Artificial Intelligence, the largest society of AI scientists and engineers, has defined artificial intelligence as the pursuit of “the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines.”²¹³ The 2019 National Defense Authorization Act includes the following definition to guide the Commission’s work:

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 an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision-making, and acting.²¹⁴

AI Has Many Uses: AI can be used for many different functions, including anomaly detection for flagging abnormal data patterns; prediction of what will happen next; recommendation of relevant alternatives; translation (between languages, for example); optimization (for example, to tune the cooling and power consumption in a data center); planning to decide what actions need to be taken to achieve a given goal; and classification by assigning an observation to a category. Different applications lead to the use of different AI approaches. This wide array of functionality allows AI applications to touch virtually every aspect of society, including the national security areas described in this report.

AI Past, Present, and Future: The term “artificial intelligence” was first coined in 1956 by a small group of computer science researchers who met at Dartmouth College.²¹⁵ They embarked on a quest to make machines “use language, form abstractions/concepts, solve problems now reserved for humans, and improve themselves.”²¹⁶ In these early years, what DARPA describes as the first wave, AI researchers explored many approaches,

including symbolic logic, machine learning, and planning. Some of the most effective results were based on “handcrafted knowledge,” where the structure of that knowledge was defined by humans and then used by the machine for reasoning and interacting.²¹⁷ These AI approaches are still very much in use, but have been augmented by a second wave of AI based on large-scale statistical machine learning that enables engineers to create models that can be trained to specific problem domains if given exemplar data (e.g., images or sensor data) or simulated interactions (e.g., game playing). These narrow machine learning (ML) approaches include genetic algorithms, probabilistic models, kernel machines, and neural networks. DARPA describes the need for a new, third wave of AI as contextual adaptation where systems will construct explanatory causal models for classes of real-world phenomena.²¹⁸ There is a growing consensus that while statistical correlation methods in the current wave of ML are both powerful and useful, they face significant limits, and new fundamental approaches will be needed.²¹⁹

Why AI Now?: Since its inception in the 1950s, AI has experienced several cycles of excitement followed by disillusionment due to unmet potential and promises. We are once again experiencing a surge of popularity, driven by innovation brought on by a convergence of factors. These include: unprecedented availability of big data; more powerful computing—particularly with the use of specialized graphics processing units (GPUs), which are well-suited for parallel computations by neural networks; ubiquitous mobile connectivity, enabling AI technologies to be easily embedded and portable while managing data within the cloud; and dramatic improvements in ML algorithms, particularly those involving deep learning (DL).²²⁰

What is Deep Learning?: DL is a statistical technique that exploits large quantities of data as training sets for a network with multiple hidden layers, called a deep neural network (DNN).²²¹ A DNN is trained on a data set, generating outputs, calculating errors, and adjusting its internal parameters. Then the process is repeated hundreds of thousands of times until the network achieves an acceptable level of performance. It has proved to be an effective technique for image classification, object detection, speech recognition, and natural language processing—problems that challenged researchers for decades. By learning from data, DNNs can solve some problems much more effectively, and also solve problems that were never solvable before.²²²

Deep Learning Has Its Challenges: While this is revolutionary, there are challenges with DNNs that can have significant implications for national security and defense applications.²²³ DNNs are data-driven, which means that unwanted or unknown biases within training data may be learned and amplified within a DNN’s decision-making. Today, DNN solutions can be influenced by small changes to input data (which appear normal to a human) that lead to unexpected results and errors. DNNs are vulnerable to data poisoning attacks, in which sometimes hard-to-spot data elements are added to the

training set in order to deliberately train the DNN to produce unwanted behavior. In DNNs, a decision is computed from a deep cascade of many parameters (weights), making it difficult for a human to understand why the decision was made. The network looks for patterns of correlation within its training set, focusing on observations of changes to its input that lead to changes to its output; however, noting a pattern of change does not provide a determined cause for the change. This is why DARPA is focusing its investments on third wave AI that will do a better job of determining causation. Any knowledge gained from DL is typically shallow and very dependent on the context represented in the training set. Reliable DNN performance in the real world presumes a largely stable and sufficiently representative sample population within the training set. This may not be the case, for example, when dealing with the changing conditions often encountered by a warfighter. These limitations must be carefully weighed when applying DL methods to any particular problem.²²⁴ They underscore the need to use it in a responsible, ethical, and risk-based manner.

Beyond Deep Learning. DNNs are typically trained using a technique known as supervised learning on massive quantities of labeled data, where the labels represent the ground truth on which the DNN is being trained. This dependence on big, labeled data poses certain challenges. The labeling process, which is usually done manually, is expensive and time-consuming. It also means DNNs only solve problems where big data is available. Many national security and defense applications require decision-making about rare events, for which only a small amount of data is available for training. There is a new programming paradigm for machine learning called “weak supervision” that does not require hand-labeled data. Instead, it uses heuristically generated training data with external knowledge bases, patterns, rules, or other classifiers.²²⁵ Reinforcement learning, where an algorithm is trained to make suitable actions by maximizing rewarded behavior over the course of its actions, is another approach.²²⁶ This type of learning can take place in simulated environments, such as game playing, which reduces the need for real-world data. One-shot (or few-shot) learning is an approach that leverages existing knowledge to enable learning in some applications (e.g., object recognition) on a few non-repeated examples, with the system rapidly learning similarities and dissimilarities between the training examples.²²⁷ AI researchers continue to explore these and other approaches that reduce the demand for real-world labeled data.

Another approach that has caught significant attention is generative adversarial networks (GANs).²²⁸ Here, two neural networks are trained in tandem: one is designed to be a generative network (the forger) and the other a discriminative network (the forgery detector). The objective is for each network to train and better itself off the other, reducing the need for big labeled training data. GANs have opened new and exciting possibilities. For example, researchers are exploring techniques for a GAN to train to normalize and generate data, in order to simulate the application domain. Once trained,

data generation and scientific experimentation can move from the real world to a digital simulation, greatly increasing the speed of experimentation and discovery. In this way, GANs present a powerful new approach to neural network training and data generation. They also pose challenges to national security by enabling “deepfakes,” as discussed elsewhere in this report.

AI Is Different from Conventional Software: The use of current data-driven machine learning makes the AI software development and management lifecycle fundamentally different from that of conventional software. The lifecycle for conventional software involves:

1. Specifying the program’s objective,
2. Implementing by writing code,
3. Testing the program, and
4. Deploying and maintaining the program.

The lifecycle for a data-driven AI system²²⁹ similarly starts with:

1. Specifying the program’s objective. In place of implementing by writing code,
2. Data is acquired and explored to meet the objective, then
3. The AI system is trained on the data.

Steps 2 and 3 are usually iterative, with data acquisition and training continuing until desired performance objectives are attained. The scale and complexity of DNNs and other data-driven ML approaches makes it more scientifically challenging to conduct the next steps:


4. Testing. Once sufficiently tested, the AI system is
5. Deployed and maintained, but this too is more complex, as there may be a
6. Data feedback loop enabling the AI system to be continuously trained and thereby improved over time even after it is deployed.

These differences in lifecycle have significant ramifications. First, normal acquisition processes are challenged. A conventional software program is typically developed by a vendor, purchased and licensed by a user, then independently deployed by the user. With AI, there is often much more interaction between the AI system’s developer and its user, particularly because of the need to access training data. There is also a growing practice of sharing pre-trained weights to support transfer learning, where further machine

learning takes place within the user's environment. These differences make the acquisition of AI technologies more like a sustainment contract rather than a traditional product purchase. Second, AI's data-driven decision-making adds significant complexity to the testing and evaluation needed to assure robustness and safety. The decision-making structure that results from DNN training is usually very complex, and as mentioned earlier, is subject to hard-to-isolate sensitivities to input data.²³⁰ Traditional approaches to software verification and validation are not sufficient for software architectures that include these structures. Testing and verifying systems that contain data-driven AI components is an area of active research. While it is strategic to support faster ML-based innovation of applications for national security, the science of testing and evaluation for AI safety must keep pace with application design in order to prevent accidents and mitigate vulnerabilities that are unique to AI systems.²³¹ If left unaddressed, shortcomings in establishing justified confidence could increase risk and hamper the rate of fielding new AI-powered systems.

Holistic View of AI: As MIT Lincoln Laboratory has explained in its AI canonical architecture,²³² AI can be holistically represented as a process pipeline. The pipeline starts with 1) the acquisition of data from sensors and sources, which is then 2) processed through data conditioning (e.g. cleaning and normalization). That creates information that can be 3) subjected to an array of possible AI approaches and machine learning techniques (e.g. knowledge-based, unsupervised and supervised learning, transfer learning, reinforcement learning, etc.). Those techniques produce knowledge, which can be 4) utilized in human-machine teaming, leading to 5) insights for mission users. Knowledge could also be provided to another machine in cases where the decision loop is too short (e.g., in cyber and hypersonic defense) and machine decision-making is required. This entire process pipeline is supported by modern computing, including conventional hardware technologies, GPUs, tensor processing units (TPUs), neuromorphic processors, custom chips, and possibly, in the future, quantum computing. Undergirding all of this is the need for the AI system to be robust. The Commission's work covers the span of this holistic view of AI.

Chart: MIT Lincoln Laboratory - AI Canonical Architecture²³³



Appendix 2: Examples of Common Challenges for DoD Contracting with Technology Companies

DoD's traditional approach to acquisition faces several common challenges that limit its attractiveness as a customer for AI companies, especially small and medium-sized businesses.

These challenges prevent the government from achieving timely access to important commercial solutions. Pockets across DoD and the IC have developed innovative, alternative approaches to procurement that begin to address these difficulties, but more remains to be done. In-Q-Tel, for example, makes strategic investments on behalf of the military and intelligence community to rapidly deliver and enhance commercial products. Operating as a non-profit, In-Q-Tel also lowers the initial and long-term costs to taxpayers by attracting \$16 of private sector investment for every \$1 of In-Q-Tel funding.²³⁴ The Defense Innovation Unit is also leveraging the Other Transaction Authority (OTA) granted by Congress to quickly contract with technology companies outside the traditional defense industrial base.²³⁵ However, barriers persist and impede the adoption of AI-enabled technologies for national security at scale.

The following four examples describe illustrative challenges DoD faces when working with tech companies:

EXAMPLE	EXPLANATION
<p>Company A assesses the DoD market and identifies opportunities but decides doing business with the government is too difficult.</p>	<p>A 2017 survey revealed slow contracting was the top complaint of startups considering working with the government.²³⁶ Startups are sprinting to find clients that will deliver cash quickly, ideally within 12 months. It is not practical for startups to pursue contracts that require two years to close, a common timeline for DoD. By then, many startups will have failed. Qualifying for DoD contracts also requires companies to stand up a separate financial accounting system to meet defense audit standards. In this example, DoD loses access to commercial solutions because of the length and complexity of the military acquisition process.</p>
<p>Company B receives a DoD pilot contract for \$750,000 that gets awarded within 60 days via Other Transaction Authority (OTA).²³⁷ After meeting its performance milestones and delivering a successful prototype in 12 months, the company’s solution never transitions to production or scales.</p>	<p>It is challenging for DoD to transition a successful prototype to a scaled production capability. If DoD deems an OTA prototype successful, it must issue a production contract before other organizations across DoD can buy the new technology. DoD’s budget cycle also lacks flexibility for adjustment in the year of execution, forcing program managers to plan years in advance. Therefore, there may be no funding available to scale even a highly successful prototype. In this example, DoD’s contracting, budgeting, and programming cycles can impede the delivery of a solution into operations.</p>
<p>Company C receives a \$750,000 contract, meets the milestones, and delivers a successful prototype. It is placed into production and the company receives \$1 million in revenue to support it. But the solution never scales beyond that.</p>	<p>Even when a single organization within DoD finds funding to place a successful prototype in production, other organizations may not be aware it is available or lack access to a common operating environment to seize its advantages. As a result, the solution never achieves its true capability through scale. In this example, DoD lacks the necessary infrastructure to share scale successes.</p>
<p>Company D hears that doing business with the government is onerous and unrewarding, so it never considers working with DoD or its investors urge against it.</p>	<p>Today, for startups seeking venture capital funding, a business plan focused on the government is almost always a deal-breaker because investors perceive accurately, based on the past two decades of experience, that DoD is a difficult customer and unrealizable market opportunity. DoD may receive \$100 billion in research, development, test, and evaluation appropriations annually,²³⁸ but this is tiny compared to the global consumer market. In this example, perceptions of DoD—and the market opportunity it represents—create selection bias among potential commercial partners, before even assessing the market.</p>

Appendix 3: AI Workforce Model

BUILD CONSENSUS			QUESTIONS FOR DEPARTMENTS	
WORKER ARCHETYPES	OUTPUT	CAPABILITIES (ETHICS THROUGHOUT)	TRAINING, EDUCATION, AND RECRUITMENT	ORGANIZATIONAL NEEDS AND COMPOSITION
AI EXPERT	Leads the ethical design, development, and deployment of AI-driven technologies; oversees test and evaluation (verification and validation) to determine technology readiness; helps maintain and leverage supporting data architecture; translates requirements into capabilities; translates technical topics for senior leaders	Expert in data science, machine learning (e.g., deep learning), AI lifecycle, applied ethics and one or more of the following: natural language processing; computer vision; robotics; human-computer interfaces; human centered systems engineering; algorithmic and computational theory	How will the national security community train or recruit and integrate AI experts?	How many AI experts does the national security workforce need? Where should they be? Should they be uniformed, civilian, or contractors?
AI DEVELOPER	Data selection and preprocessing; model selection, training, and validation; partnership with domain knowledge experts and end users; discovery of local opportunities	Computational statistics and data science; programming (e.g. Python or R); model development using an ML library	Who trains developers for the national security workforce? When will they be identified and trained?	How many developers does the national security workforce need? Should they be uniformed, civilian, contractors, and/or contracted companies?
DEPLOYMENT SPECIALIST	Infrastructure installation and maintenance, review input/output sent by end-users, additions to training data sets, rough examination of training data sets, training/testing existing models, deployment	Hardware/software installation and maintenance, training data management, model verification/validation, algorithm deployment, data cleansing	Education equivalent to a technical certification offered by a military program or vocational training	How many AI technicians does the national security workforce need? Where should they be?
END USER	Daily business augmented/enabled by AI	Use of systems and apps	Normal systems training	Ubiquitous
NON-TECHNICAL TACTICAL LEADER	Gathers tactical requirements to guide the development of new AI-enabled capabilities, oversees deployment to ensure tactical requirements are met; partners with technicians, data engineers, and AI experts; leads normal operations	Tactical domain implementation expert, basic data collection and management, basic understanding of AI decision making within the context of use and the sources of failures and errors, ethics applied to tactical use	How will the national security community train and educate tactical leaders? How much do they need to know?	How many tactical leaders should the national security enterprise have?
NON-TECHNICAL STRATEGIC LEADER	Oversees the creation of strategic and enterprise objectives, considers the ethics of new capabilities, oversees deployment and scaling, partnership with experts, developers, and tactical leaders; career management	Basics and ethics of AI lifecycle, strategic and enterprise expertise, tactical domain management, software development processes	When and where will leaders learn about AI? How much do they need to know?	How will leaders incentivize AI competence? How many leaders need to be competent, and at what point in their careers?
SUPPORT ROLES	Acquisition and contracting of AI hardware and software, services, and identification of commercial opportunities; legal support; legislative affairs, human resources, etc.	Understanding of software purchasing, data boundaries/limitations and rights; funding requirements; compute purchases, identification of skill and qualifications of AI practitioners; legal and ethical aspects of development and deployment	When and where will support experts learn about AI? How much do they need to know?	What parts of the support workforce needs to learn about AI demands?

Table: AI Workforce Model²³⁹

The AI Workforce Model was developed by the NSCAI in partnership with the Defense Innovation Board and the Joint Artificial Intelligence Center. Their collaboration on the model does not extend to the remainder of the report. The material is based on more than 30 briefings with experts from AI-first companies, traditional companies that have successfully integrated AI, consulting groups, AI organizations within the government, and human resource and force structure experts within the government. The model and explanatory note also include information from AI and organizational theory discussed in business and academic literature.

Adopting a common workforce model will help DoD approach AI workforce development with a common set of concepts, vocabulary, and understanding of the types of questions it needs to answer. This model describes different types of AI workers, their outputs, skills, training and education, and composition and disposition in the larger workforce. The model is meant to serve as a tool to guide DoD's understanding of workforce needs, not as a set of recommendations for career fields.

The most important takeaway from this model is that building an AI workforce will require much more than highly educated, deep technical experts. DoD must also develop non-technical leaders, deployment specialists, and end-users to effectively employ AI solutions across the force.

Ideally, the first three columns will be endorsed by consensus throughout DoD and the U.S. government:

- *Column 1: Worker Archetypes.* The model has seven worker archetypes that should be represented in an AI workforce.
- *Column 2: Output.* The output column describes what each category of worker will contribute.
- *Column 3: Capabilities.* The capabilities column lists critical, required knowledge and abilities.

The last two columns offer guiding questions for which each department and agency will likely have different answers based on different enterprise strategies and needs:

- *Column 4: Training/Education/Recruitment.* The education/training/recruitment column asks how the government will develop each type of worker.

- *Column 5: Organizational needs and composition:* The far-right column, organizational needs and composition, asks how many of each type of worker the government will need, where they will be located within departments and agencies, and what percentage of them will be uniformed, civilian, or contractors that work alongside government counterparts, or contracted companies that deliver a service.

AI WORKER ARCHETYPES

Below, the archetypes shown in the model graphic are explained in more detail, including a non-exhaustive list of sub-archetypes with an example persona.

Technical Roles

AI Experts will lead the ethical design, development, and deployment of AI-driven technologies; translate requirements into capabilities; and help inform senior leaders. The greatest difference between AI developers and AI experts will be experts' ability to oversee testing and evaluation, an area that AI developer training may not support adequately enough to sufficiently minimize risk. AI experts are expected to have the educational, work, and research experience equivalent to a PhD.

- Sub-archetypes: AI research engineer, AI software and systems architect, AI machine learning software engineer, cloud computing application architect, solution architect, machine learning engineer, human-centered systems engineer.
- Example job illustration:
 - AI research engineers focus on research and development of technologies that enable and advance semi- and fully-autonomous systems. They serve as algorithm experts with up-to-date knowledge of modern AI research and may be involved in the inception of ideas and drive the development cycles from research to testing of prototypes for a major project or component of a major project.
 - AI solution architects identify and collect data sources, analyze and extract key data and information, and evaluate and monitor data quality to meet the organization's information system needs and requirements.

AI Developers will be data focused. They will be responsible for data cleaning, feature extraction and selection, and analysis; model training and tuning; partnerships with domain knowledge experts and end users; and the discovery of local opportunities for exploitation. Developers require less training and education than AI experts, and will have training, education, and/or experience that is roughly equivalent to an associate or bachelor's degree; and that includes relevant ethics and bias mitigation in data

processing and model training. Because they require less training than experts, there is more potential for the government to hire or internally train developers and to have them more widespread across the workforce. This allows the U.S. government to expand the pool of AI talent it selects and trains, placing less reliance on a small number of universities and private sector companies whose relationship with the national security enterprise may not always be relied upon.

- Sub-archetypes: data engineer, data analyst, data administrator, software engineer.
- Example job illustration:
 - Data engineers deliver full-stack data solutions across the entire data processing pipeline and rely on systems engineering principles to implement solutions that span the data lifecycle to collect, ingest, process, store, persist, access, and deliver data at scale and at speed. They have knowledge of local, distributed, and cloud-based technologies; data virtualization and smart caching; and all security and authentication mechanisms required to protect data.

Deployment Specialists will be responsible for the installation and maintenance of the hardware/software that collects and processes data, the regular management of end user inputs and outputs, and management of data sets. They will be the most common point of contact with technical expertise for end users. They are strongly analogous to today's mechanics and IT specialists and technicians.

- Sub-archetypes: AI hardware engineer, AI systems engineer.
- Example job illustration:
 - AI hardware/software engineers serve as hardware/software experts for autonomous systems and work with other experts to provide the next generation of hardware/software solutions, including and not limited to sensing computer storage as well as controls and systems safety. They support teams with the integration of hardware with software and systems, human-machine interface tests, and preparations of autonomous systems for certification and deployment.

Enablers

End users will use AI-enabled systems during normal operations. Their use of AI will strongly resemble the use of currently available software in that it will require some system-specific training, but, with the exception of some positions that manage data, little to no AI specific expertise. Most if not all members of the federal government will be end users.

- Sub-archetypes: tracked vehicle mechanic, all-source intelligence analyst, F-35 pilot.

Non-technical tactical leaders will serve as domain knowledge experts that help create the tactical requirements for AI systems, ensure the effective and ethical employment of AI systems, and partner with developers and experts. Tactical leaders already exist in today's organization, such as the military's officer corps. To become part of an effective AI workforce, they should be trained to understand the basics of data collection and management, AI decision making, and AI specific ethics.

- Sub-archetypes: Battalion/squadron commander, program manager, senior intelligence analyst.

Non-technical strategic leaders will oversee the creation of strategic and enterprise objectives, the deployment and scaling of new systems, and manage the careers of developers and experts. All organizations need to train and certify their strategic leaders in areas such as the basics and ethics of the AI lifecycle and software development processes, in order to be able to interpret and trust output from AI-enabled decision support systems.

- Sub-archetypes: Deputy Assistant Secretary of Defense, U.S. Central Command Commander, J7, Deputy Director of National Intelligence for Mission Integration.

Support roles is the broadest category on the workforce model, and includes the support functions that are necessary to support AI development and employment. These include, but are not limited to acquisition officers who understand how to identify and purchase viable or modifiable commercial solutions and contracting officers who can negotiate service and development contracts that address traditionally troubled topics like data rights; human resource officers who understand how to leverage hiring authorities to quickly and less painfully hire talented developers and experts and the skills and qualifications of AI practitioners; legislative affairs personnel need to be able to explain AI funding requirements to members of Congress and their staff; and legal professionals need to understand the legal and ethical aspects of the entire AI development and deployment process.

- Sub-archetypes: human resource specialist (classification/recruitment & placement), legislative fellow, staff judge advocate.

Appendix 4: Organizations Consulted

From March 2019 to October 2019, the Commission's staff engaged with a wide variety of entities that play a role in AI and national security issues. Based on the Commission's mandate from Congress and in an effort to gain a diverse set of opinions, Commission staff received briefings from stakeholders in government, industry, academia, non-profits, associations, and individual experts from around the United States, and from international partners. This appendix alphabetically lists many of the Commission staff's engagements. The Commission staff will continue its engagements through the remainder of 2019 and into 2020 before issuing the Final Report. If your organization is interested in contacting the Commission, feel free to email the Commission at inquiry@nscai.gov. (*Note: An organization's inclusion on this list does not mean endorsement of the Commission's Interim Report.*)

Aerospace Industry Association	Defense Advanced Research Projects Agency
AFWERX	Defense Innovation Board
AI Sustainable Development Group	Defense Innovation Unit
Air Force Research Lab	Defense Intelligence Agency
Algorithmic Warfare Cross Functional Team, DoD	Deloitte
Amazon	Department of Commerce
American Psychological Association	Department of Defense
Anduril	Department of Energy
Arizona State University	Department of Homeland Security
Army Futures Command	Department of State
Army Research Lab	Department of the Air Force
Army War College	Department of the Army
Asia America Multi Technology Association	Department of the Navy
Association for the Advancement of AI	Department of the Treasury
Atlantic Council	Draper Laboratory
Australian Strategic Policy Institute	Duke University
Battery Innovation Center	Electronic Frontier Foundation
Booz Allen Hamilton	Elsevier
Brookings Institution	Energy Systems Network
Bureau of Industry & Security	Ethical Intelligence Consulting
C3IOT	Eurasia Group
California Polytechnic State University	Federal Bureau of Investigation
Carnegie Endowment for International Peace	Federation of American Scientists
Carnegie Mellon University	Franklin Templeton Investments
Center for a New American Security	Future of Privacy Forum
Center for Democracy & Technology	General Dynamics
Center for Naval Analysis	Georgetown University
Central Intelligence Agency	Google
Cisco Systems	Govini
Coding it Forward	Harvard University
Computer Science & Telecommunications Board	Harvard-MIT Ethics & Governance of AI Initiative
Computing Research Association	Heritage Foundation
Cybersecurity & Infrastructure Security Agency	Howard University
Cyberspace Solarium Commission	Human Rights Watch for the Campaign to Stop Killer Robots
Data & Society	IBM

Intelligence Advanced Research Projects Activity
 In-Q-Tel
 Indiana Economic Development Corporation
 Indiana General Assembly
 Indiana Innovation Institute
 Indiana Office of Defense Innovation
 Indiana University
 Institute for Defense Analyses
 International Committee of the Red Cross
 International Embassies
 Johns Hopkins University
 John Hopkins University Applied Physics Lab
 Joint Artificial Intelligence Center
 The Joint Staff, DoD
 Kessel Run
 Lockheed Martin
 Marine Corps
 Marine Corps Warfighting Laboratory
 Massachusetts Institute of Technology
 McKinsey
 Microsoft
 MIT Lincoln Laboratory
 National Commission on Military, National & Public Service
 National Defense University
 National Geospatial-Intelligence Agency
 National Institute of Standards & Technology
 National Reconnaissance Office
 National Science Foundation
 National Security Agency
 National Security Council
 National Security Innovation Network
 Naval Sea Systems Command
 Naval Surface Warfare Center - Crane
 Naval Undersea Warfare Center Division – Keyport
 Navy Digital Warfare Office
 Networking & Information Technology Research & Development Program
 New York University
 Northrop Grumman
 Notre Dame University
 NVIDIA
 Odium Strategies, LLC
 Office of Civil Liberties, Privacy & Transparency. ODNI
 Office for Civil Rights & Civil Liberties, DHS
 Office of Commercial & Economic Analysis, DoD
 Office of the Director of National Intelligence
 Office of Management and Budget
 Office of Naval Research
 Office of Personnel Management
 Office of the Secretary of Defense
 Office of Science & Technology, DHS
 Office of Science & Technology Policy, The White House
 OpenAI
 Pacific Northwest National Lab
 Palantir Technologies
 Partnership for Public Service
 Paulson Institute
 Presidential Innovation Fellows
 Privacy Office, DHS
 Primer.ai
 Privacy & Civil Liberties Oversight Board
 Purdue University
 Radius Indiana
 RAND
 Raytheon
 Reagan Institute
 SAP National Security Services
 Schmidt Futures
 Semiconductor Industry Association
 SensorHound
 Shield AI
 SIMBA Chain
 Singularity University
 SoftBank
 Software Engineering Institute
 SOSI
 Stanford University
 Tech Inquiry
 The Engine
 The Technical Cooperation Partnership
 Tufts University
 U.S. House of Representatives
 U.S. International Trade Commission
 U.S. Senate
 U.S. Special Operations Command
 U.S.-China Economic & Security Review Commission
 United States Air Force Academy
 United States Military Academy
 United States Naval Academy
 United States Naval War College
 University of California System
 University of California, Berkeley
 University of Chicago
 University of Illinois at Urbana-Champaign
 University of Oxford
 University of Pennsylvania
 University of Southern California
 University of Southern Indiana
 University of Washington
 University of Washington Applied Physics Lab
 Yale University

Appendix 5: Commission Staff and Advisors

Yll Bajraktari, *Executive Director*

Michael Gable, *Chief of Staff*

Commission Staff:

Seth Center
Matt Cordova
Caroline Danauy
Tess deBlanc-Knowles
Rama Elluru
Michael Garris
Chelsea Holt
Charles Howell
Mike Jackson
Rebekah Kennel
Jeffrey Kojac
Lance Lantier
Michael Lueptow
Justin Lynch
Paul Maykish
Kevin McGinnis
Christopher McGuire
Brandon McKee
Brent Myles
Robert Nelson
Angela Ponmakha
Paul Rhodes
Christopher Rice
Tara Rigler
Jennifer Sheehan

Jenilee Keefe Singer

Claire Trotter

Jackson Valen

Zoe Weinberg

Kate Yeager

Jessica Young

Olivia Zetter

Advisors to the Commission:

John Bansemer

Susanna Blume

David Danks

Jeff Ding

Kathleen Fisher

Michael Horowitz

Elsa Kania

Christopher Kirchhoff

Frank Long

Brendan McCord

Heather Roff


Nadia Schadow Murphy

Paul Sharre

William Sherlis

Raj Shah

Amy Zegart



Endnotes

¹ Pub. L. 115-232, The John S. McCain National Defense Authorization Act for Fiscal Year 2019, 132 Stat. 1636, 1964 (2018). [hereinafter FY 2019 NDAA]

² China's principal AI strategy, issued in 2017, envisions that "by 2030, China's AI theories, technologies, and applications should achieve world-leading levels, making China the world's primary AI innovation center." For a full translation, see Graham Webster, Rogier Creemers, Paul Triolo, and Elsa Kania, *Full Translation: China's New Generation Artificial Intelligence Development Plan*, New America (Aug. 1, 2017), <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.

³ Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press at 22-23 (2014).

⁴ Dave Martinez, Nick Malyska, Bill Streilein et al., *Artificial Intelligence: Short History, Present Developments, and Future Outlook*, MIT Lincoln Laboratory at 27 (Jan. 2019), <https://www.ll.mit.edu/media/9526>. [hereinafter MIT Lincoln Laboratory AI Study Report]

⁵ Note that model-based AI requires data for the manual construction of the model(s). Typically, this involves less data than statistical machine learning, but more human effort.

⁶ Saleema Amershi et al., *Software Engineering for Machine Learning: A Case Study*, CSE-SEIP '10 Proceedings of the 41st International Conference on Software Engineering at 291-300 (2019), <https://2019.icse-conferences.org/details/icse-2019-Software-Engineering-in-Practice/30/Software-Engineering-for-Machine-Learning-A-Case-Study>; D. Sculley et al., *Machine Learning: The High Interest Credit Card of Technical Debt*, <https://ai.google/research/pubs/pub43146>.

⁷ Counter-AI broadly includes security measures within AI systems to protect them from subversion by bad actors and efforts to deter and defend against the malicious use of AI.

⁸ For an overview, see Nils J. Nilsson, *The Quest for Artificial Intelligence*, Cambridge University Press (2010).

⁹ See e.g., Michael Horowitz, *Artificial Intelligence, International Cooperation, and the Balance of Power*, Texas National Security Review (May 2018), <https://tnsr.org/2018/05/artificial-intelligence-international-competition-and-the-balance-of-power/>.

¹⁰ Effectively programming AI applications for anomaly detection requires a careful design process that accounts for nuance within the indicator set. This is particularly important in counterterrorism and policing use cases, which need to account for the fact that not all anomalies are hostile actions. See Emily Berman, *A Government of Laws and Not of Machines*, Boston University Law Review at 1322-23 (2018), <https://www.bu.edu/bulawreview/files/2018/10/BERMAN.pdf>.

¹¹ See Amy Zegart and Michael Morrell, *Spies, Lies, and Algorithms: Why US Intelligence Agencies Must Adopt or Fail*, Foreign Affairs (May/June 2019), <https://www.foreignaffairs.com/articles/2019-04-16/spies-lies-and-algorithms>; Amy Zegart, *U.S. Intelligence Needs Another Reinvention*, The Atlantic (Sept. 11, 2019),

<https://www.theatlantic.com/ideas/archive/2019/09/us-intelligence-needs-another-reinvention/597787/>.

¹² The Defense Science Board Summer Study on Autonomy, which inspired further DoD efforts to develop autonomous systems, determined that autonomy is fueled by advances in AI and identified two categories of intelligent systems: those employing autonomy at rest (operating virtually) and those employing autonomy in motion (operating in the physical world). The study found that autonomy delivers significant military value and that DoD must accelerate its adoption of autonomous capabilities. *Summer Study on Autonomy*, Defense Science Board, Department of Defense at 1, 5 (June 9, 2016), <https://dsb.cto.mil/reports/2010s/DSBSS15.pdf>. Current DoD policy on autonomy in weapon systems is governed by a DoD Directive. Directive No. 3000.09, Autonomy in Weapon Systems, Department of Defense (Nov. 21, 2012, updated May 8, 2017), <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>.

¹³ For example, DoD established an Algorithmic Warfare Cross-Functional Team in 2017—also known as Project Maven—in order to “accelerate DoD’s integration of big data and machine learning” and “turn the enormous volume of data available to DoD into actionable intelligence and insights at speed.” Memorandum from the Deputy Secretary of Defense (Apr. 26, 2017), https://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_awcft_project_maven.pdf.

¹⁴ NATO was established 70 years ago, in 1949. In Asia, it has been almost 70 years since deterrence failed to forestall major war on the Korean Peninsula in 1950.

¹⁵ There is disagreement regarding the extent to which AI could identify second-strike capabilities, potentially enabling an adversary to launch a successful decapitation strike in a nuclear stand-off. Such a capability would jeopardize the United States’ survivability of nuclear forces. See Zachary S. Davis, *Artificial Intelligence on the Battlefield: An Initial Survey of Potential Implications for Deterrence, Stability, and Strategic Surprise*, Center for Global Security Research, Lawrence Livermore National Laboratory (Mar. 2019), https://cgsr.llnl.gov/content/assets/docs/CGSR-AI_BattlefieldWEB.pdf; Edward Geist and Andrew J. Lohn, *How Might Artificial Intelligence Affect the Risk of Nuclear War?*, RAND Corporation (2018) <https://www.rand.org/pubs/perspectives/PE296.html>; Vincent Boulain, Ed., *The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk*, Vol. I, Euro-Atlantic Perspectives, Stockholm International Peace Research Institute (May 2019), <https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf>; Zachary Kallenborn, *AI Risks to Nuclear Deterrence are Real*, War on the Rocks (Oct. 10, 2019), <https://warontherocks.com/2019/10/ai-risks-to-nuclear-deterrence-are-real/>; and Raphael Loss and Joseph Johnson, *Will Artificial Intelligence Imperil Nuclear Deterrence?*, War on the Rocks (Sept. 19, 2019), <https://warontherocks.com/2019/09/will-artificial-intelligence-imperil-nuclear-deterrence/>.

¹⁶ See, e.g., Miles Brundage et al., *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation* (Feb. 2018), <https://arxiv.org/pdf/1802.07228.pdf>. For one view on possible terrorist uses of AI, see Jacob Ware, *Terrorist Groups, Artificial Intelligence, and Killer Drones*, War on the Rocks (Sept. 24, 2019), <https://warontherocks.com/2019/09/terrorist-groups-artificial-intelligence-and-killer-drones/>.

¹⁷ See, e.g., Matt Chessen, *The MADCOM Future: How Artificial Intelligence Will Enhance Computational Propaganda, Reprogram Human Culture, and Threaten Democracy...And What Can Be Done About It*, Atlantic Council (2017), https://www.atlanticcouncil.org/wp-content/uploads/2017/09/The_MADCOM_Future_RW_0926.pdf; Robert Chesney and Danielle Citron, *Deepfakes and the New Disinformation War*, Foreign Affairs (Jan.-Feb. 2019),

<https://www.foreignaffairs.com/articles/world/2018-12-11/deepfakes-and-new-disinformation-war>; Miles Brundage et al., *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation* (Feb. 2018). Pew recently found that “[a]bout two-thirds of Americans (66%) say they at least sometimes come across altered videos and images that are intended to mislead, with 15% encountering them often.” See Jeffrey Gottfried, *About Three-Quarters of Americans Favor Steps to Restrict Altered Videos and Images*, Pew Research Center (Jun. 24, 2019), <https://www.pewresearch.org/fact-tank/2019/06/14/about-three-quarters-of-americans-favor-steps-to-restrict-altered-videos-and-images/>. On the potential impact of Deepfakes on news, see, e.g., Asa Fitch, *Readers Beware: AI Has Learned to Create Fake News Stories*, The Wall Street Journal (Oct. 13, 2019), <https://www.wsj.com/articles/readers-beware-ai-has-learned-to-create-fake-news-stories-11571018640>. AI’s ability to automate, accelerate, and scale synthetic accounts and content on social media can “hyperpower Russia’s use of disinformation.” Alina Polyakova, *Weapons of the Weak: Russia and AI-Driven Asymmetric Warfare*, Brookings (Nov. 15, 2018), <https://www.brookings.edu/research/weapons-of-the-weak-russia-and-ai-driven-asymmetric-warfare/>.

¹⁸ Steven Feldstein, *The Global Expansion of AI Surveillance*, Carnegie Endowment for International Peace (Sept. 2019), https://carnegieendowment.org/files/WP-Feldstein-AISurveillance_final1.pdf. [hereinafter Feldstein, AI Surveillance Paper]

¹⁹ See, e.g., William Dixon and Nicole Eagan, *3 Ways AI will Change the Nature of Cyber Attacks*, World Economic Forum (June 19, 2019), <https://www.weforum.org/agenda/2019/06/ai-is-powering-a-new-generation-of-cyberattack-its-also-our-best-defence/>; Dustin Frazee, *Cyber Grand Challenge (CGC)*, Defense Advanced Research Projects Agency, <https://www.darpa.mil/program/cyber-grand-challenge>.

²⁰ Particularly alarming is evidence of Russian propaganda and malicious content targeting active service members and veterans on Facebook, Twitter, and other social media platforms. Kristofer Goldsmith, *An Investigation Into Foreign Entities Who are Targeting Troops and Veterans Online*, Vietnam Veterans of America (Sept. 2019), <http://vva.org/trollreport/>; see also Marcus Comiter, *Attacking Artificial Intelligence: AI’s Security Vulnerability and What Policymakers Can Do About It*, Harvard Belfer Center (Aug. 2019), <https://www.belfercenter.org/sites/default/files/2019-08/AttackingAI/AttackingAI.pdf>; and Jim Baker, *Artificial Intelligence - A Counterintelligence Perspective: Part 3*, Lawfare (Oct. 10, 2018), <https://www.lawfareblog.com/counterintelligence-implications-artificial-intelligence-part-iii>.

²¹ AI systems reacting to one another in a national security context could generate unintended consequences that are “destructive or counterproductive,” and may occur “before human users can adequately respond.” Paul Scharre and Michael Horowitz, *Artificial Intelligence: What Every Policymaker Needs to Know*, Center for a New American Security (June 2018), <https://www.cnas.org/publications/reports/artificial-intelligence-what-every-policymaker-needs-to-know>; and Richard Danzig, *Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority*, Center for a New American Security (June 2018), <https://www.cnas.org/publications/reports/technology-roulette>. [hereinafter Danzig, Technology Roulette]

²² Reports indicate that a “cohort of countries is moving toward digital authoritarianism by embracing the Chinese model of extensive censorship and automated surveillance systems.” Adrian Shahbaz, *The Rise of Digital Authoritarianism*, Freedom House (Oct. 2018), <https://freedomhouse.org/report/freedom-net/freedom-net-2018/rise-digital-authoritarianism>.

²³ On blacklists and “social credit” systems see, e.g., Jamie Horsley, *China’s Orwellian Social Credit Score Isn’t Real*, Foreign Policy (Nov. 16, 2018), <https://foreignpolicy.com/2018/11/16/chinas-orwellian-social-credit-score-isnt-real/>; Louise Matsakis, *How the West Got China’s Social Credit System Wrong*, WIRED (July

29, 2019), <https://www.wired.com/story/china-social-credit-score-system/>; Shazeda Ahmed, *The Messy Truth About Social Credit*, Logic Magazine (May 1, 2019), <https://logicmag.io/china/the-messy-truth-about-social-credit/>; and Shazeda Ahmed, *Credit Cities and the Limits of the Social Credit System*, in Shazeda Ahmed et al., *AI, China, Russia, and the Global Order: Technological, Political, Global, and Creative Perspectives*, Department of Defense (Dec. 2018), https://nsiteam.com/social/wp-content/uploads/2018/12/AI-China-Russia-Global-WP_FINAL.pdf.

²⁴ Human Rights Watch has documented the Chinese government's use of mass surveillance systems to this end: "Xinjiang authorities conduct compulsory mass collection of biometric data, such as voice samples and DNA, and use artificial intelligence and big data to identify, profile, and track everyone in Xinjiang." Maya Wang, *Eradicating Ideological Viruses: China's Campaign of Repression Against Xinjiang's Muslims*, Human Rights Watch (Sept. 9, 2018), <https://www.hrw.org/report/2018/09/09/eradicating-ideological-viruses/chinas-campaign-repression-against-xinjiangs>.

²⁵ Public reports indicate that some of Russia's robotic combat systems have not performed well under real-world conditions, but they indicate a willingness to experiment, learn, and adapt. See Dr. Margarita Konaev and Samuel Bendett, *Russian AI-Enabled Combat: Coming to a City Near You?*, War on the Rocks (July 31, 2019), <https://warontherocks.com/2019/07/russian-ai-enabled-combat-coming-to-a-city-near-you/>.

²⁶ *2018 Global R&D Funding Forecast*, R&D Magazine at 3 (Winter 2018).

²⁷ Field Cady and Oren Etzioni, *China to Overtake US in AI Research*, Allen Institute AI2Blog (Mar. 13, 2019), <https://medium.com/ai2-blog/china-to-overtake-us-in-ai-research-8b6b1fe30595>.

²⁸ Tencent and Alibaba are both multi-hundred billion-dollar companies. Bytedance, the Beijing based maker of TikTok, the fastest growing social media app in the world, is the most valuable private tech company in the world with a \$78 billion valuation. Yingzhi Yang & Julie Zhu, *TikTok Owner ByteDance's H1 Revenue Better Than Expected at Over \$7 Billion*, Reuters (Sept. 30, 2019), <https://www.reuters.com/article/us-bytedance-results-exclusives/exclusive-tiktok-owner-bytedances-h1-revenue-better-than-expected-at-over-7-billion-sources-idUSKBN1WF0G7>.

²⁹ Members of Congress have raised concerns regarding TikTok's data collection practices and its possible national security ramifications. See Letter from U.S. Senators Charles E. Schumer & Tom Cotton to Acting Director of National Intelligence Joseph Maguire (Oct. 23, 2019), <https://www.democrats.senate.gov/imo/media/doc/10232019%20TikTok%20Letter%20-%20FINAL%20PDF.pdf>.

³⁰ See Meng Jing, *China to Boost its 'National Team' to Meet Goal of Global AI Leadership by 2030*, South China Morning Post (Nov. 15, 2018), <https://www.scmp.com/tech/innovation/article/2173345/china-boost-its-national-team-meet-goal-global-ai-leadership-2030>.

³¹ On military-civil fusion, see Testimony of Elsa B. Kania before the U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion: China's Pursuit of Artificial Intelligence*, New Materials and New Energy (June 7, 2019), https://www.uscc.gov/sites/default/files/June%202019%20Hearing_Panel%201_Elsa%20Kania_Chinese%20Military%20Innovation%20in%20Artificial%20Intelligence_0.pdf; and Lorand Laskai, *Civil-Military Fusion: The Missing Link Between China's Technological and Military Rise*, Council on Foreign Relations (Jan. 29, 2018), <https://www.cfr.org/blog/civil-military-fusion-missing-link-between-chinas-technological-and-military-rise>.

³² Russia established the Advanced Research Foundation (ARF) in 2012. In 2018, ARF announced new proposals for the Ministry of Defense focused on image and speech recognition and autonomous military systems. Samuel Bendett, *The Rise of Russia's Hi-Tech Military*, American Foreign Policy Council (June 26, 2019). The Chinese Central Military Commission has established a Science and Technology Commission, as well as a Military Scientific Research Committee, which will oversee the PLA's future research efforts and is intended to advance military and technological innovation. Elsa B. Kania, *Battlefield Singularity: Artificial Intelligence, Military Revolution, and China's Future Military Power*, Center for a New American Security (Nov. 2017), <https://www.cnas.org/publications/reports/battlefield-singularity-artificial-intelligence-military-revolution-and-chinas-future-military-power>.

³³ Reports indicate that Chinese military strategists anticipate an evolution from today's "informatized" warfare to a new era of "intelligentized" warfare. Elsa B. Kania, *Chinese Military Innovation in Artificial Intelligence*, Center for a New American Security (June 7, 2019), <https://www.cnas.org/publications/congressional-testimony/chinese-military-innovation-in-artificial-intelligence>.

³⁴ A recent study found that "[f]or the foreseeable future, U.S. demand for AI talent will far outpace domestic supply." Zachary Arnold et al., *Immigration Policy and the U.S. AI Sector*, Georgetown Center for Security and Emerging Technology at 1 (Sept. 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_Immigration_Policy_and_AI.pdf.

³⁵ "China's talent recruitment plans, such as the Thousand Talents Program, offer competitive salaries, state-of-the-art research facilities, and honorific titles, luring both Chinese overseas talent and foreign experts alike to bring their knowledge and experience to China, even if that means stealing proprietary information or violating export controls to do so." Testimony of Bill Priestap, Assistant Director for Counterintelligence, Federal Bureau of Investigation, before the U.S. Senate Judiciary Committee, *Hearing on China's Non-Traditional Espionage Against the United States* at 5 (Dec. 12, 2018), <https://www.judiciary.senate.gov/imo/media/doc/12-12-18%20Priestap%20Testimony.pdf>.

³⁶ "Other countries, including U.S. allies and China, strategically respond to changes in U.S. immigration policy, and they have significantly increased efforts to recruit U.S.-based AI talent." Zachary Arnold et al., *Immigration Policy and the U.S. AI Sector*, Georgetown Center for Security and Emerging Technology at 6 (Sept. 2019). For example, Canada's Global Skills Strategy has lured several thousand from the United States since it started in 2017, according to the Canadian Immigration Minister. See Theophilos Argitis, *Canada Is Luring Tech Talent Away from U.S. with Fast-Track Visa*, Bloomberg (June 12, 2019), <https://www.bloomberg.com/news/articles/2019-06-12/canada-is-luring-tech-talent-away-from-u-s-with-fast-track-visa>. At the same time, recent years have shown a decline in the number of new international students attending American colleges, from 300,743 in 2015 to 271,738 in 2017. See Institute for International Education, *International Student Enrollment Trends 1948/49-2017/28, Open Doors Report on International Student Exchange* (2018), <http://www.iie.org/opendoors>.

³⁷ There has been a robust dialogue concerning the comparative progress of the United States and China on AI, with varying perspectives on the current state of the race and its future trajectory. Our preliminary judgments are based on a broad-based assessment of the data and trends that the Commission has surveyed to date. A comprehensive review of the literature and methodology on the relative capabilities of China and the United States is beyond the scope of this Interim Report; for more detailed assessments, see Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order*, Houghton Mifflin Harcourt (2018); Gregory C. Allen, *Understanding China's AI Strategy*, Center for New American Security (Feb. 6, 2019), <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>; Yoav Shoham,

Raymond Perrault, Erik Brynjolfsson, Jack Clark et al., *The AI Index 2018 Annual Report*, AI Index Steering Committee, Human Centered AI Initiative, Stanford University (Dec. 2018), <http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf>; Daniel Castro, Michael McLaughlin, and Eline Chivot, *Who is Winning the AI Race: China, the EU, or the United States?*, Center for Data Innovation (Aug. 2019), <http://www2.datainnovation.org/2019-china-eu-us-ai.pdf>; *China AI Development Report*, China Institute for Science and Technology Policy, Tsinghua University (July 2018); *Artificial Intelligence: How Knowledge is Created, Transferred, and Used, Trends in China, Europe, and the United States*, Elsevier (Dec. 2018), <https://www.elsevier.com/connect/resource-center/artificial-intelligence>; Field Cady and Oren Etzioni, *China May Overtake US in AI Research*, Allen Institute for Artificial Intelligence (Mar. 13, 2019), <https://medium.com/ai2-blog/china-to-overtake-us-in-ai-research-8b6b1fe30595>; and *China's AI Dream*, MacroPolo, <https://macropolo.org/digital-projects/chinai/chinai-intro/>.

³⁸ One analysis found that of authors at U.S. institutions with papers accepted to the Conference on Neural Information Processing Systems (NeurIPS), about 9 percent were likely Chinese nationals, based on the location of their undergraduate degree. Joy Dantong Ma, *The AI Race is Wide Open, If America Remains Open*, MacroPolo (Apr. 15, 2019), <https://macropolo.org/us-china-ai-race-talent/>. One survey indicates that 14 percent of Silicon Valley technical professionals are of Chinese origin. *Share of Residents in Technical Occupations with a Bachelor's Degree or Higher, by Place of Origin*, Silicon Valley Institute for Regional Studies (2017), <https://siliconvalleyindicators.org/data/people/talent-flows-diversity/tech-talent/share-of-residents-in-technical-occupations-with-a-bachelors-degree-or-higher-by-place-of-origin/>.

³⁹ *A Conversation with Christopher Wray*, Council on Foreign Relations (Apr. 26, 2019), <https://www.cfr.org/event/conversation-christopher-wray-0>.

⁴⁰ See Testimony of John C. Demers, Assistant Attorney General, National Security Division, U.S. Department of Justice, before the U.S. Senate Committee on the Judiciary, *Hearing on China's Non-traditional Espionage against the United States: The Threat and Potential Policy Responses* (Dec. 12, 2018), https://www.justice.gov/sites/default/files/testimonies/witnesses/attachments/2018/12/18/12-05-2018_john_c_demers_testimony_re_china_non-traditional_espionage_against_the_united_states_the_threat_and_potential_policy_responses.pdf.

⁴¹ See, e.g., *Review of the FY2020 Budget Request for NIH*, Hearing of the Subcommittee on Labor, Health and Human Services, Education, and Related Agencies, Senate Committee on Appropriations (Apr. 11, 2019), <https://www.appropriations.senate.gov/hearings/review-of-the-fy2020-budget-request-for-nih>; William C. Hannas, James Mulvenon & Anna B. Puglisi, *Chinese Industrial Espionage: Technology Acquisition and Military Modernization*, Routledge at 89, 168, 171-4 (2013).

⁴² See *Review of the FY2020 Budget Request for NIH*, Hearing of the Subcommittee on Labor, Health and Human Services, Education, and Related Agencies, Senate Committee on Appropriations (Apr. 11, 2019); Jeffrey Mervis, *NIH Probe of Foreign Ties Has Led to Undisclosed Firings—and Refunds from Institutions*, Science Magazine (June. 26, 2019), <https://www.sciencemag.org/news/2019/06/nih-probe-foreign-ties-has-led-undisclosed-firings-and-refunds-institutions>.

⁴³ Alex Joske, *Picking Flowers, Making Honey: The Chinese Military's Collaboration with Foreign Universities*, Australian Strategic Policy Institute (Oct. 30, 2018), <https://www.aspi.org.au/report/picking-flowers-making-honey>.

⁴⁴ For example, reports indicate that while the imposition of strict export controls would decrease the two countries' interdependence, "fully disentangling the global supply chain would impose enormous

economic costs.” Chris Meserole, *Artificial Intelligence and the Security Dilemma*, Brookings (Nov. 6, 2018), <https://www.brookings.edu/blog/order-from-chaos/2018/11/06/artificial-intelligence-and-the-security-dilemma/>.

⁴⁵ Ranked by number of faculty publications accepted at the top AI conferences, U.S. universities occupied three of the top five rankings, and six of the top ten in 2018. Matt Sheehan, Joy Dantong Ma, and Jeffrey Ding, *The Talent*, MacroPolo, 2019, <https://macropolo.org/chinai/the-talent/>. A report by MIT’s Lincoln Laboratory holds that the top eight universities in AI are located in the United States. MIT Lincoln Laboratory AI Study Report at 8. And according to a report by Tsinghua University, U.S. universities occupied seven of the top ten universities for top AI talent. *China AI Development Report*, China Institute for Science and Technology Policy, Tsinghua University at 37 (July 2018).

⁴⁶ Remco Zwetsloot, Roxanne Heston, and Zachary Arnold, *Strengthening the U.S. AI Workforce: A Policy and Research Agenda*, Center for Security and Emerging Technology at iii (Sep. 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_U.S._AI_Workforce.pdf; see also Science & Engineering Indicators 2018, National Science Board (2018), <https://www.nsf.gov/statistics/2018/nsb20181/assets/901/tables/tt03-27.pdf>.

⁴⁷ See *AI 100: the Artificial Intelligence Startups Redefining Industries*, CBInsights (2019), <https://www.cbinsights.com/research/artificial-intelligence-top-startups/>.

⁴⁸ Of the 32 companies classified by CBInsights as AI Unicorns, 17 are American. *The Increasingly Crowded AI Unicorn Club*, CBInsights (Feb. 14, 2019), <https://app.cbinsights.com/research/ai-unicorn-club/>.

⁴⁹ See Testimony of Helen Toner, Director of Strategy, Center for Security and Emerging Technology, Georgetown University, before the U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion: China’s Pursuit of Artificial Intelligence, New Materials, and New Energy* (June 7, 2019), <https://cset.georgetown.edu/wp-content/uploads/Toner-USCC-Testimony-6.7.19.pdf>.

⁵⁰ The Honorable Donald J. Trump, *Executive Order on Maintaining American Leadership in Artificial Intelligence*, The White House (Feb. 11, 2019), <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>. [hereinafter Executive Order on AI]

⁵¹ See *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*, Department of Defense (2018), <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF>. [hereinafter DoD AI Strategy Summary]

⁵² See *The AIM Initiative: Augmenting Intelligence Using Machines Increasing Insight and Knowledge through Artificial Intelligence, Automation and Augmentation*, Office of the Director of National Intelligence (2019), www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf. [hereinafter ODNI AIM Initiative]

⁵³ See Department of Energy’s Artificial Intelligence and Technology Office website, <https://www.energy.gov/science-innovation/artificial-intelligence-and-technology-office>.

⁵⁴ The National Science Foundation, for example, was established in 1950. On the triangular alliance, see, e.g., Walter Isaacson, *The Sources of America’s Innovative Edge*, in Joseph Nye et al., *Technology and National Security: Maintaining America’s Edge*, Aspen Institute (2019).

⁵⁵ Notably, the combined annual R&D investment of only the top five U.S. tech firms exceeds that of the DoD. See *2018 Global R&D Funding Forecast*, R&D Magazine at 17 (Winter 2018).

⁵⁶ Federal R&D has dropped from a height of 1.86 percent of GDP in 1964 to just 0.61 percent in 2017. See *Federal R&D Budget Dashboard*, American Association for the Advancement of Science, <https://www.aaas.org/programs/r-d-budget-and-policy/federal-rd-budget-dashboard>.

⁵⁷ Academic-corporate collaborations account for nearly nine percent of published AI research in the United States over the past twenty years. *Artificial Intelligence: How Knowledge is Created, Transferred, and Used*, Elsevier at 57 (2018).

⁵⁸ The U.S. government continues to fund about half of university-based scientific R&D. See “Chapter 5 Highlights” of *Science and Engineering Indicators 2018*, National Science Foundation (2018), <https://nsf.gov/statistics/2018/nsb20181/report/sections/academic-research-and-development/highlights>.

⁵⁹ Due to the method of reporting in prior years, there is no definitive data on prior-year government spending on AI R&D. The Federal Networking and Information Technology Research and Development (NITRD) Program measures R&D investments along designated Program Component Areas (PCAs). 2019 is the first year in which AI is designated as a PCA. See *The Networking and Information Technology Research & Development Program Supplement to the President’s FY2020 Budget*, Executive Office of the President at 4 (Sep. 2019), <https://www.whitehouse.gov/wp-content/uploads/2019/09/FY2020-NITRD-AI-RD-Budget-September-2019.pdf>. [hereinafter R&D Program Supplement to FY2020 Budget]

⁶⁰ *Federal R&D Obligations Increase an Estimated 2.7 percent in FY 2018*, National Science Foundation (2018), <https://www.nsf.gov/statistics/2019/nsf19321/#fig2>.

⁶¹ *Analysis of Current and Future Computer Science Needs via Advertised Faculty Searches for 2019*, CRA Bulletin (Dec. 7, 2019), <https://cra.org/analysis-of-current-and-future-computer-science-needs-via-advertised-faculty-searches-for-2019/>.

⁶² Significant data gaps make it difficult to assess with precision the level of China’s public investment in AI R&D. While some reports have suggested that China is spending tens of billions of dollars, other analyses indicate that the level of investment in 2018 was closer to a few billion dollars. See, e.g., Thomas Colvin et al., *A Tentative Framework for Examining U.S. and Chinese Expenditures for Research and Development on Artificial Intelligence*, Institute for Defense Analyses (Sep. 2019).

⁶³ In 1953, the U.S. spent 0.72 percent of its GDP on R&D. In 1957, when the then-Soviet Union launched Sputnik, it had grown to 1.3 percent. R&D spending peaked at 1.86 percent in 1964. In 2017, it declined below 1953 levels to 0.61 percent. *Federal R&D Budget Dashboard*, American Association for the Advancement of Science, <https://www.aaas.org/programs/r-d-budget-and-policy/federal-rd-budget-dashboard>.

⁶⁴ *FY2020 White House budget request for Research and Development*, The White House at 271, https://www.whitehouse.gov/wp-content/uploads/2019/03/ap_21_research-fy2020.pdf.

⁶⁵ See *2018 Global R&D Funding Forecast*, R&D Magazine at 5 (Winter 2018).

⁶⁶ See Tony Peng and Michael Sarazen, *Are Commercial Labs Stealing Academia's AI's Thunder?*, SyncedReview (July 10, 2019), <https://medium.com/syncedreview/are-commercial-labs-stealing-academias-ai-thunder-dd51cf4bd8d6>. One study identified 180 AI faculty from North American universities who departed academia for an industry job from 2004-2018. In a sign of the acceleration of the trend, 40 of these moves occurred in 2018 alone. See Michael Gofman and Zhao Jin, *Artificial Intelligence, Human Capital, and Innovation*, University of Rochester (Aug. 20, 2019), http://gofman.info/AI/AI_GofmanZhao.pdf.

⁶⁷ These measures do not include spending of the defense or intelligence community on AI R&D, which is not reported publicly. See *R&D Program Supplement to FY 2020 Budget* at 10.

⁶⁸ The most qualified proposals include those rated “competitive” or “highly competitive” in the peer review process. In 2018, NSF funded \$165 million in core AI research, but did not have room in the budget to fund another \$185 million worth of well-rated proposals. In 2017, the numbers were \$122 million funded, \$174 million unfunded. NSF presentation to NSCAI (June 2019).

⁶⁹ *AI Next Campaign*, DARPA, <https://www.darpa.mil/work-with-us/ai-next-campaign>.

⁷⁰ Peter Highnam, *Artificial Intelligence Colloquium: DARPA Future R&D in AI*, YouTube (Mar. 29, 2019), <https://www.youtube.com/watch?v=tl-Yfl27ijU>.

⁷¹ *National Artificial Intelligence (AI) Research Institutes*, National Science Foundation, https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505686.

⁷² The National Nanotechnology Initiative (NNI) was launched in 2000 to coordinate 25 departments and agencies. The initiative spurred a “thriving interdisciplinary nanotechnology community of about 150,000 contributors,” a “flexible R&D infrastructure consisting of about 100 large nanotechnology-oriented R&D centers, networks, and user facilities,” and an “expanding industrial base of about 3,000 companies producing nanotechnology-enabled products.” Mihail C. Roco, Chad A. Mirkin, and Mark C. Hersam, *Nanotechnology Research Directions for Societal Needs in 2020: Retrospective and Outlook*, *Springer Science & Business Media* (June 17, 2011). We thank a research team at the University of Pennsylvania for pointing us to the NNI model. University of Pennsylvania Input to the National Security Commission on Artificial Intelligence, Working Paper (Sep. 12, 2019). [hereinafter U. Penn. Working Paper] For more information see National Nanotechnology Initiative, <https://www.nano.gov/>.

⁷³ Y. Gil and B. Selman, *A 20-Year Community Roadmap for Artificial Intelligence Research in the US*, Computing Community Consortium and the Association for the Advancement of Artificial Intelligence at 95 (Aug. 2019), <https://cra.org/ccc/wp-content/uploads/sites/2/2019/08/Community-Roadmap-for-AI-Research.pdf>

⁷⁴ See *What are Clusters?*, Harvard Business School's Institute for Strategy & Competitiveness, <https://www.isc.hbs.edu/competitiveness-economic-development/frameworks-and-key-concepts/Pages/clusters.aspx>; Mark Muro, *Countering the Geographical Impacts of Automation: Computers, AI, and Place Disparities*, The Brookings Institution (Feb. 14, 2019), <https://www.brookings.edu/research/countering-the-geographical-impacts-of-automation-computers-ai-and-place-disparities/>.

⁷⁵ Established as part of Canada’s Pan-Canadian AI Strategy. See *Pan-Canadian Artificial Intelligence Strategy*, Invest in Canada, <https://www.investcanada.ca/why-invest/pan-canadian-artificial-intelligence-strategy>.

⁷⁶ The Tech Bridges program has created five regional hubs across the country: in Newport, Rhode Island; Keyport, Washington; San Diego, California; Orlando, Florida; and Crane, Indiana. “These Tech Bridges will partner with start-ups, academia, corporations, small businesses, non-profits, private capital, and government entities.” *Tech Bridges*, NavalX, U.S. Department of the Navy, <https://www.secnv.navy.mil/agility/Pages/techbridges.aspx>.

⁷⁷ The Naval Surface Warfare Center (NSWC) in Crane, Indiana, has created a tech hub with substantial engineering expertise (103 PhDs on staff). A recent MIT study highlighted NSWC’s ability to serve DoD’s engineering needs while also acting as a catalyst for economic development. NSWC has partnered with top regional universities (Purdue, Notre Dame, and Indiana University), as well as corporations, non-profits, and state government. The result is a strong track record of support for the Navy’s strategic systems, electronic warfare, and expeditionary warfare—while spurring growth in surrounding rural areas. See Kathryn Person, Dylan Cohen, Jonathan Miller, and Fiona Murray, *NSWC Crane Innovation Analysis: Contributing to Regional Innovation Ecosystems*, MIT Innovation Initiative (May 17, 2019), [https://innovation.mit.edu/assets/NSWC-Crane-Report .pdf](https://innovation.mit.edu/assets/NSWC-Crane-Report.pdf).

⁷⁸ For a related idea, see the DoD AI Strategy, which proposes “seeding new AI innovation districts.” DoD AI Strategy Summary at 12.

⁷⁹ Currently, Howard Hughes Medical Institute (HHMI) supports nearly 300 individual investigators. Twenty-nine current or former HHMI-affiliated investigators have received a Nobel Prize. U. Penn. Working Paper. For more information, see <https://www.hhmi.org/>.

⁸⁰ The Vannevar Bush fellowships provide five-year, single-investigator awards of up to \$3 million. See *Vannevar Bush Faculty Fellowship*, Office of the Under Secretary of Defense for Research & Engineering, <https://basicresearch.defense.gov/Programs/Vannevar-Bush-Faculty-Fellowship/>.

⁸¹ AI research requires, at a minimum, access to capital-intensive computation, specialized data sets, and data storage; for larger-scale ambitions, engineering support from software and hardware engineers is also needed.

⁸² NSF is piloting a cloud access entity to connect NSF-funded researchers with cloud computing resources. See CloudBank, <https://www.cloudbank.org/>. Such a model could be expanded to provide cloud access to a broader portion of the research community.

⁸³ The government could unlock further advancements by spurring the development of physical test beds that enable researchers to apply AI to energy, transportation, agriculture, manufacturing, and other capital-intensive industries.

⁸⁴ Executive Order on AI.

⁸⁵ *The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update*, Select Committee on Artificial Intelligence of the National Science and Technology Council (June 2019), <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>.

⁸⁶ See R&D Program Supplement to FY2020 Budget.

⁸⁷ The details of the NSPM are not public, but it is referenced in the *Executive Order on Maintaining American Leadership in Artificial Intelligence*.

⁸⁸ Paul Scharre has argued that “[o]ne of the most important ways policymakers can deal with the dangers of AI is to boost funding for AI safety research” and advocated for the creation of “red teams” within DoD and DHS that test systems through simulated attacks. Paul Scharre, *Killer Apps: The Real Dangers of an AI Arms Race*, Foreign Affairs (May/June 2019), <https://www.foreignaffairs.com/articles/2019-04-16/killer-apps>.

⁸⁹ Among the few exceptions are IARPA’s TrojAI program and DARPA’s GARD program. See also, Nathan Strout, *The Three Major Security Threats to AI*, C4ISRNET (Sept. 10, 2019), <https://www.c4isrnet.com/artificial-intelligence/2019/09/10/the-3-major-security-threats-to-ai/>.

⁹⁰ Congress has taken steps to examine the threat posed by deepfakes. For example, the House Intelligence Committee held hearings on deepfakes and synthetic media in June 2019. In addition, the Deepfake Report Act of 2019 (S.2065, H.R.3600), introduced in July 2019 and passed by the Senate in October, reflects a bipartisan effort to address the potential misuse of AI technology by malicious actors, including its use by foreign governments to harm national security.

⁹¹ For a helpful overview of important differences between the needs of the commercial and national security sectors in applying AI, see MIT Lincoln Laboratory AI Study Report at 29. A more detailed look into DoD’s priorities for AI research can be found at AI Next Campaign, DARPA, <https://www.darpa.mil/work-with-us/ai-next-campaign>.

⁹² Important FFRDCs for AI-related work include MIT’s Lincoln Laboratory and Carnegie Mellon’s Software Engineering Institute, among others. FFRDCs close to Silicon Valley, such as the Livermore and Berkeley National Laboratories, could also serve as hubs connecting leading AI researchers to national missions.

⁹³ *Defense Science and Technology: Actions Needed to Enhance Use of Laboratory Initiated Research Authority*, Government Accountability Office (Dec. 2018), <https://www.gao.gov/assets/700/696192.pdf>.

⁹⁴ See, e.g., Robert O. Work and Greg Grant, *Beating the Americans at their Own Game: An Offset Strategy with Chinese Characteristics*, Center for a New American Security (June 2019), <https://s3.amazonaws.com/files.cnas.org/documents/CNAS-Report-Work-Offset-final-B.pdf?mtime=20190531090041>.

⁹⁵ See Robert Martinage, *Toward a New Offset Strategy: Exploiting U.S. Long-Term Advantages to Restore U.S. Global Power Projection Capability*, Center for Strategic and Budgetary Assessments (Oct. 2014), <https://csbaonline.org/uploads/documents/Offset-Strategy-Web.pdf>.

⁹⁶ Memorandum from the Secretary of Defense, to the Deputy Secretary of Defense et al. (Nov. 15, 2014), <https://archive.defense.gov/pubs/OSD013411-14.pdf>; Remarks by Robert O. Work, Deputy Secretary of Defense, delivered at the Reagan Defense Forum, *The Third Offset Strategy* (Nov. 7, 2015), <https://www.defense.gov/Newsroom/Speeches/Speech/Article/628246/reagan-defense-forum-the-third-offset-strategy/>.

⁹⁷ The Commission recognizes that the advent of AI has broader implications across several other national security departments and agencies. Our initial focus, however, has been on applications within DoD and the IC.

⁹⁸ Michael C. Horowitz, *The Diffusion of Military Innovation: Causes and Consequences for International Politics*, Princeton University Press (2010).

⁹⁹ For a similar point, see Kelley Sayler, *Artificial Intelligence and National Security*, Congressional Research Service (Jan. 30, 2019).

¹⁰⁰ See, e.g., Sally Cole, *Cognitive Electronic Warfare: Countering Threats Posed by Adaptive Radars*, Military Embedded Systems, <http://mil-embedded.com/articles/cognitive-electronic-warfare-countering-threats-posed-by-adaptive-radars/>.

¹⁰¹ See, e.g., ACTUV “Sea Hunter” Prototype Transitions to Office of Naval Research for Further Development, DARPA (Jan. 30, 2018), <https://www.darpa.mil/news-events/2018-01-30a>.

¹⁰² *National Security Strategy of the United States of America*, The White House (Dec. 2017), <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>; *Summary of the 2018 National Defense Strategy*, Department of Defense (2018), <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>; *DoD Digital Modernization*, Department of Defense (July 12, 2019), <https://media.defense.gov/2019/Jul/12/2002156622/-1/-1/1/DOD-DIGITAL-MODERNIZATION-STRATEGY-2019.PDF>; ODNI AIM Initiative.

¹⁰³ Testimony of the Honorable Mark T. Esper, Secretary of Defense, before the Senate Armed Services Committee at 64 (July 16, 2019), https://www.armed-services.senate.gov/imo/media/doc/19-59_07-16-19.pdf.

¹⁰⁴ Brendan McCord, Eye on AI Transcript, (Aug. 28, 2019), <https://static1.squarespace.com/static/5b75ac0285ede1b470f58ae2/t/5d6aa8edb91b0c0001c7a05f/1567271149944/EP+22+brendan+transcript.pdf>.

¹⁰⁵ See, e.g., *Acquisition Authority*, U.S. Special Operations Command, <https://www.socom.mil/SOF-ATL/Pages/Acquisition-Authority.aspx>.

¹⁰⁶ Joint Artificial Intelligence Center, Office of the Chief Information Officer, Department of Defense, <https://www.ai.mil/index.html>.

¹⁰⁷ Testimony of Dana Deasy, Chief Information Officer, Department of Defense, before the House Armed Services Committee, Subcommittee on Emerging Threats and Capabilities at 6 (Dec. 11, 2018), <https://docs.house.gov/meetings/AS/AS26/20181211/108795/HHRG-115-AS26-Wstate-DeasyD-20181211.pdf>.

¹⁰⁸ ODNI AIM Initiative.

¹⁰⁹ *Id.*

¹¹⁰ See J. Michael McQuade et al., *Acquisition of Software-Defined Hardware-Based Adaptable Systems*, Center for Strategic and International Studies (Aug. 2019); *Final Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulation*, Section 809 Panel, Vol. 3 (Jan. 2019), <https://section809panel.org/volume-3-report/>; and J. Michael McQuade et al., *Software Is Never Done: Refactoring the Acquisition Code For Competitive Advantage*, Defense Innovation Board (May 2019), https://media.defense.gov/2019/Mar/26/2002105909/-1/-1/0/SWAP.REPORT_MAIN.BODY.3.21.19.PDF.

¹¹¹ Implementation of the DIB software study recommendations is led by the DoD Under Secretary for Acquisition and Sustainment. Related legislative reforms are included in the FY 2020 NDAA, such as requiring the Secretary of Defense to issue guidance for “rapid acquisitions” of software; and “software development and software acquisition training and management programs.” H.R. 2500 & S. 1790, National Defense Authorization Act for Fiscal Year 2020, Title VIII, Subtitle A, §§ 801-802.

¹¹² The Air Force-MIT AI Accelerator was established in 2019 for “rapid prototyping, scaling, and application of AI algorithms and systems,” and has a \$75 million budget over five years. See Rob Matheson, *MIT and U.S. Air Force Sign Agreement to Launch AI Accelerator*, MIT News (May 20, 2019), <http://news.mit.edu/2019/mit-and-us-air-force-sign-agreement-new-ai-accelerator-0520>. Challenge teams work on Air Force problems, and the program utilizes MIT Lincoln Laboratory as an intermediary to translate classified projects into the unclassified domain and vice versa. *Id.* Kessel Run is another Air Force partnership with MIT. Kessel Run has its own acquisition authority and can therefore spend money as needed to move technology from R&D to fielding. See *Kessel Run*, <https://kesselrun.af.mil/>. The Army has set up an AI Task Force at Carnegie Mellon. See Gary Sheftik, *AI Task Force Taking Giant Leaps Forward*, U.S. Army (Aug. 13, 2019), https://www.army.mil/article/225642/ai_task_force_taking_giant_leaps_forward.

¹¹³ See DoD FY 2020 Budget Estimates: Office of the Secretary of Defense, Defense-wide Justification Book, Volume 3 of 5, Research, Development, Test & Evaluation, Defense-wide at 1-171 to 1-173 https://comptroller.defense.gov/Portals/45/Documents/defbudget/fy2020/budget_justification/pdfs/03_RDT_and_E/RDTE_Vol3_OSD_RDTE_PB20_Justification_Book.pdf.

¹¹⁴ The AI vulnerability surface encompasses the entire AI training environment, including data and models, and, upon deployment, a wide range of system connections (e.g., sensors, outputs, and network connections). All of these must be secured individually and collectively. For the military, AI applications are always embedded in a larger system-of-systems, and the security of the entire system must be addressed to ensure reliable operations and resilience.

¹¹⁵ See *DoD Digital Modernization Strategy*, Department of Defense (July 12, 2019); *DoD AI Strategy Summary*; and *The United States Air Force Artificial Intelligence Annex to the Department of Defense Artificial Intelligence Strategy*, Department of the Air Force (2019), <https://media.defense.gov/2019/Sep/12/2002182176/-1/-1/1/US%20AIR%20FORCE%20AI%20ANNEX%20TO%20DOD%20AI%20STRATEGY.PDF>.

¹¹⁶ NSCAI interview with U.S. government officials (May 24, 2019).

¹¹⁷ Four points about methodology:

1) Our research on AI workforce issues has focused mainly on DoD so far, both civilian personnel and the uniformed military. Other parts of the “national security AI workforce” include the intelligence community, DHS, DOE, State, and the national labs and FFRDCs that undertake work for

defense and intelligence missions. Much of the PhD-level AI expertise resides in these labs and research institutes.

2) It is difficult to collect reliable quantitative data about the federal AI workforce because of complications in identifying types of practitioners based on job title or category. Our assessment so far relies heavily on more than 50 interviews conducted between May and August 2019.

3) It is also hard to identify qualified AI workers based on academic credentials alone. Carnegie Mellon offers a specific degree in AI, but other schools that consistently have top-ranked AI programs do not. People with academic training in AI most commonly study computer science, mathematics, physics, neuroscience, psychology, and philosophy. But only a fraction of the people who study those subjects graduate with a background in AI. That may change over time as AI becomes a more popular specialty in academic programs.

4) We recognize there is some overlap in talking about an AI workforce and a broader “cyber” or “digital” workforce. But in the defense and intelligence contexts, there are differences between cyber operations and AI. While the skills demanded of cyber warriors and AI experts overlap, they are not the same, and they have somewhat different roles and missions. Cyber espionage and warfare involve “digital maneuver” to achieve some aim. AI is a technical discipline able to create capabilities in every operating and military domain, through some sort of computation. AI is focused on creating algorithms that might be used in the cyber domain, but can also be used in non-networked hardware. Hacking is also a different skill than data engineering and model training.

¹¹⁸ As Richard Danzig has observed: “Limited understanding hobbles policymakers’ discussions and decisions about whether to develop, how to employ, and how much to rely on complex capabilities. Increasingly, senior officials are called to make decisions about, and on the basis of, technologies that did not exist at the time of their education.” Danzig, *Technology Roulette*.

¹¹⁹ More AI-focused military education is a good place to start. The service academies and professional military education institutes are beginning to incorporate AI into their curricula. As they develop custom programs, they can leverage existing resources in the meantime—including AI boot camps at universities and in the private sector—to improve leader education now. For example, the Hoover Institution at Stanford ran a bootcamp in August 2019 geared for congressional staff. See *Cyber and Artificial Intelligence Boot Camp*, Hoover Institution (Aug. 2019), <https://www.hoover.org/events/cyber-and-artificial-intelligence-boot-camp-2019>.

¹²⁰ The Defense Digital Service has created such an environment, but DDS is small and exceptional. See *Defense Digital Service*, Department of Defense, <https://dds.mil/>.

¹²¹ Isaac Porche, Caoliann O’Connell, and John Davis, *Cyber Power Potential of the Army’s Reserve Component*, RAND Corporation at xiii (2017).

¹²² These could include online courses. For a sampling of online courses in AI, see Bernard Marr, *The 6 Best Free Online Artificial Intelligence Courses Today*, *Forbes* (Apr. 16, 2018), <https://www.forbes.com/sites/bernardmarr/2018/04/16/the-6-best-free-online-artificial-intelligence-courses-for-2018/#dbb1d7759d75>.

¹²³ NSCAI interview with U.S. Government official (June 25, 2019).

¹²⁴ Steve Hirsch, *USAF Looking for Airmen Who Speak Computer*, Air Force Magazine (Sept. 11, 2018), <http://airforcemag.com/Features/Pages/2018/September%202018/USAF-Looking-for-Airmen-Who-Speak-Computer.aspx>.

¹²⁵ A program like the Cyber Excepted Service for DoD can allow agencies to increase flexibility in hiring. See *DOD Cyber Excepted Service Personnel System*, Chief Information Officer, Department of Defense, <https://dodcio.defense.gov/Cyber-Workforce/CES.aspx>. Notably, it is time consuming for agencies to apply for excepted service status, and acceptance is competitive.

¹²⁶ For DoD, the law allows up to 2,500 HQEs to serve for up to five years, with another year at the Secretary's discretion. See 5 U.S.C. § 9903.

¹²⁷ One IC agency employs more than 400 interns annually and has an 80 percent acceptance rate for job offers. The internship program is so useful that an agency official told us they would double the number of students if their budget allowed. NSCAI staff interview (June 24, 2019).

¹²⁸ The program piloted a STEM-specific “track” for the classes of 2014-2016, which was subsequently eliminated in 2017. We are examining how the program can best be utilized for AI-related hiring. See *2017 Assessment Preparation Guide*, Presidential Management Fellows (Nov. 14, 2016), https://www.pmf.gov/media/99807/2017_pmf_assess_prep_guide_11-14-16.pdf.

¹²⁹ DoD's SMART scholarship has brought STEM talent into the Department for over a decade. *SMART Scholarship Program*, Department of Defense, <https://smartscholarshipprod.service-now.com/smart>. The Cyber Corps (Scholarship for Service), a recruiting program managed by NSF, OPM, and DHS, could be utilized to attract AI talent. Recipients are obligated to work in an agency for a period of time equal to the time covered by the scholarship. About 3,600 students have participated since 2001, training at 70 academic institutions. *History / Overview, CyberCorps: Scholarship for Service*, Office of Personnel Management, <https://www.sfs.opm.gov/Overview-History.aspx>.

¹³⁰ See, e.g., the Secretary of Defense Executive Fellows program, <https://prhome.defense.gov/Readiness/EducationTraining/SDEF.aspx>.

¹³¹ For example, our early impression is that the DoD Cyber Information Technology Exchange Program could be more fully utilized. See *Cyber Information Technology Exchange Program*, Department of Defense, <https://dodcio.defense.gov/In-the-News/Information-Technology-Exchange-Program/>. On obstacles, see, e.g., C. Todd Lopez, *DOD to Take Over Background Checks by Fiscal 2020*, Department of Defense (June 25, 2019), <https://www.defense.gov/explore/story/Article/1886923/dod-to-take-over-background-checks-by-fiscal-2020/>; J. Michael McQuade et al., *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, Department of Defense at S72, S155 (May 3, 2019).

¹³² See, e.g., the Cyber Security Exchange Act, which was introduced in the Senate in February 2019. Cyber Security Exchange Act, S.429, 116th Cong. (2019).

¹³³ An IPA agreement is essentially a detail or secondment; depending on the terms, the agency or the home institution can cover the costs. The IPA is also a useful way for agencies to draw on experts from FFRDCs. According to OPM, “agencies do not take full advantage of the IPA program.” See *Hiring Information: Intergovernment Personnel Act*, Office of Personnel Management, <https://www.opm.gov/policy-data-oversight/hiring-information/intergovernment-personnel-act/>.

¹³⁴ Cade Metz, *Tech Giants are Paying Huge Salaries for Scarce A.I. Talent*, New York Times (Oct. 22, 2017), <https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html>; see also Remco Zwetsloot et al., *Strengthening the U.S. AI Workforce*, Georgetown Center for Security and Emerging Technology (Sept. 2019).

¹³⁵ As NSA's General Counsel recently wrote: "A large portion of the intelligence community's experts on the military capabilities and plans of Russia and China joined government during the Reagan administration; other experts on counterterrorism and new technology burnished their technical skills following the Sept. 11 attacks. Many of those experts are nearing retirement or have already left to join an attractive private sector." Glenn Gerstell, *I Work for the NSA. We Cannot Afford to Lose the Digital Revolution*, New York Times (Sept. 10, 2019), <https://www.nytimes.com/2019/09/10/opinion/nsa-privacy.html>.

¹³⁶ *Id.* As a Stanford undergraduate explained: "One of the main reasons people pick companies is they want to do social good. . . . People would laugh if the government said the only way to be impactful is to work in government." Amy Zegart & Kevin Childs, *The Divide Between Silicon Valley and Washington is a National Security Threat*, The Atlantic (Dec. 13, 2018), <https://www.theatlantic.com/ideas/archive/2018/12/growing-gulf-between-silicon-valley-and-washington/577963/>.

¹³⁷ Similar to the White House Fellows program, such a program would select the most talented students for what some have envisioned as "a prestigious, one-year, high-impact stint" in government, working directly for senior defense leaders. *Id.*

¹³⁸ James Manyika & William H. McRaven, *Innovation and National Security*, Council on Foreign Relations at 66 (2019), <https://www.cfr.org/report/keeping-our-edge/>. [hereinafter CFR Report: Innovation & National Security]

¹³⁹ "Collectively, the limited evidence suggests that AI, as a field, is even less diverse than computer science as a whole." Meredith Whittaker et al., *AI Now Report 2018*, New York University (Dec. 2018), https://ainowinstitute.org/AI_Now_2018_Report.pdf. The Council on Foreign Relations Task Force reported that "[m]inorities and women remain underrepresented in STEM fields. Only 2.2 percent of Latinos, 2.7 percent of African Americans, and 3.3 percent of American Indians and Alaska Natives hold a university degree in STEM fields. Women constitute 47 percent of the overall workforce but only 28 percent of the science and engineering workforce, and women in tech jobs leave the field at a rate 45 percent higher than men." CFR Report: Innovation & National Security at 26.

¹⁴⁰ "A more diverse AI community will be better equipped to anticipate, spot, and review issues of unfair bias and better able to engage communities likely affected by bias." Jake Silberg & James Manyika, *Tackling Bias in Artificial Intelligence (and in Humans)*, McKinsey & Company (June 2019), <https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans>.

¹⁴¹ *The Current Landscape of Computer Science Enrollments in Assessing and Responding to the Growth of Computer Science Undergraduate Enrollments*, National Academies of Sciences, Engineering, and Medicine (2018), <https://www.nap.edu/read/24926/chapter/5>.

¹⁴² CRA Enrollment Committee Institution Subgroup, *CS Undergraduate Enrollments Surge Since 2006*, Computer Research Association (2017), <https://cra.org/data/Generation-CS/>.

¹⁴³ Paul Basken, *Engineering and Computer Science: Time to Separate?*, Prism (Sept. 2018), <http://www.asee-prism.org/engineering-and-computer-science-time-to-separate/>. The *Wall Street Journal* reported that “[t]he share of newly minted U.S. computer-science Ph.D.s taking industry jobs has risen to 57% from 38% over the last decade, according to data from the National Science Foundation. Though the number of Ph.D.s in the field has grown, the proportion staying in academia has hit a historic low.” Daniela Hernandez & Rachael King, *Universities’ AI Talent Poached by Tech Giants*, Wall Street Journal (Nov. 24, 2016), <https://www.wsj.com/articles/universities-ai-talent-poached-by-tech-giants-1479999601>.

¹⁴⁴ CFR Report: Innovation & National Security at 41.

¹⁴⁵ Georgia Tech’s online M.S. program in computer science is one example. We thank a team of researchers at the University of Pennsylvania for bringing this to our attention. The University of Pennsylvania is also pursuing scalable, online AI training programs. U. Penn. Working Paper.

¹⁴⁶ Testimony of Jeff Ding before the U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion: China’s Pursuit of Artificial Intelligence, New Materials and New Energy* at 3 (June 7, 2019), https://www.uscc.gov/sites/default/files/June%207%20Hearing_Panel%20Jeffrey%20Ding_China%27s%20Current%20Capabilities%2C%20Policies%2C%20and%20Industrial%20Ecosystem%20in%20AI.pdf.

¹⁴⁷ Remco Zwetsloot et al., *Strengthening the U.S. AI Workforce*, Center for Security and Emerging Technology at 4 (Sept. 2019). A study by researchers at the University of Pennsylvania pointed out that almost 60 percent of all degrees in AI in the United States are awarded to non-resident aliens. The head of the Allen Institute has noted that two-thirds of the Institute’s research scientists are immigrants. U. Penn. Working Paper.

¹⁴⁸ Jean-Francois Gagné et al., *Global AI Talent Report 2019*, JF Gagne (Apr. 2019), <https://jfgagne.ai/talent-2019>.

¹⁴⁹ *Id.* However, a recent Council on Foreign Relations Task Force report argued that “U.S. failure to compete for global talent could result in the loss of its lead” in AI. CFR Report: Innovation & National Security at 47.

¹⁵⁰ Georgetown’s Center for Security and Emerging Technology has developed proposals on AI and U.S. immigration policy that we will examine along with other perspectives. See Zachary Arnold et al., *Immigration Policy and the U.S. AI Sector*, Center for Security and Emerging Technology (Sept. 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_Immigration_Policy_and_AI.pdf. The foreign composition of the U.S. AI talent pool also raises questions for the government and military—related, for example, to U.S. citizenship requirements and security clearances. As we explore this issue, we expect that restrictions on the ability of non-citizens to participate directly in government work will likely underscore the importance of building commercial and academic partnerships to fully leverage the talent available in American society.

¹⁵¹ For example, arXiv.org collects scientific papers for public use.

¹⁵² For a summary of China’s tech transfer activities, see Sean O’Connor, *How Chinese Companies Facilitate Technology Transfer from the United States*, U.S.-China Economic and Security Review Commission (May

2019), <https://www.uscc.gov/sites/default/files/Research/How%20Chinese%20Companies%20Facilitate%20Tech%20Transfer%20from%20the%20US.pdf>. While cyber theft and industrial espionage have received a great deal of attention, the primary practices by which China has executed AI-related tech transfer are through legal mechanisms, for example, via direct technology purchases, cooperative ventures, investments, and education of Chinese nationals abroad. See Wm. C. Hannas & Huey-meei Chang, *China's Access to Foreign AI Technology*, Center for Security and Emerging Technology (Sept. 2019), <https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-Access-to-Foreign-AI-Technology-2.pdf>.

¹⁵³ Testimony of The Honorable Daniel R. Coats, Director of National Intelligence, before the U.S. Senate Select Committee on Intelligence, *Hearing on Worldwide Threat Assessment* at 14 (Jan., 29, 2019), <https://www.dni.gov/files/ODNI/documents/2019-ATA-SFR---SSCI.pdf>.

¹⁵⁴ See, e.g., Remarks by The Honorable Christopher Wray, Director of the FBI, delivered at the Council on Foreign Relations, *The FBI and the National Security Landscape* (Apr. 26, 2019), <https://www.fbi.gov/news/speeches/the-fbi-and-the-national-security-threat-landscape-the-next-paradigm-shift>. While espionage is one threat, universities also face other indirect challenges, as noted earlier in this report.

¹⁵⁵ The semiconductor industry manufactures computer chips necessary for training and execution (also called “inference”) of AI systems. The AI chip market is expected to grow faster than and take an increasingly greater share of the wider semiconductor industry. This AI-driven demand for computer chips is changing the nature of semiconductor R&D, with less focus on general-purpose computer chips like central processing units (CPUs) and more focus on more specialized computer chips that exhibit high speed and energy efficiency for AI systems. These chips include graphics processing units (GPUs), which are most commonly used for training; field-programmable gate arrays (FPGAs), which are most commonly used for inference; and application-specific integrated circuits (ASICs), which are used for both. For further background see, e.g., Gaurav Batra et al., *Artificial-Intelligence Hardware: New Opportunities for Semiconductor Companies*, McKinsey & Company (Jan. 2019), <https://www.mckinsey.com/industries/semiconductors/our-insights/artificial-intelligence-hardware-new-opportunities-for-semiconductor-companies>; Neil Thompson & Svenja Spanuth, *The Decline of Computers as a General Purpose Technology: Why Deep Learning and the End of Moore's Law are Fragmenting Computing*, available at SSRN (Nov. 20, 2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3287769.

¹⁵⁶ Michaela D. Platzer & John F. Sargent Jr., *U.S. Semiconductor Manufacturing: Industry Trends, Global Competition, Federal Policy*, Congressional Research Service (June 2016).

¹⁵⁷ SME includes the equipment needed to produce the silicon, create the electrical circuits, test the equipment, and assemble it.

¹⁵⁸ These countries have even more significant advantages in advanced SME equipment, such as photolithography tools, which are almost exclusively made in the Netherlands and Japan.

¹⁵⁹ For current controls, see *Commerce Control List: Category 3 - Electronics*, Bureau of Industry and Security (May 23, 2019), <https://www.bis.doc.gov/index.php/documents/regulations-docs/federal-register-notice/federal-register-2014/990-ccl3-2/file>.

¹⁶⁰ See *Review of Controls for Certain Emerging Technologies*, Dept. of Commerce, 83 FR 58201 (Nov. 19, 2018), <https://www.federalregister.gov/documents/2018/11/19/2018-25221/review-of-controls-for-certain-emerging-technologies>. This proposed rule emerged from the Export Control Reform Act of 2018, which was part of the FY 2019 NDAA.

¹⁶¹ For instance, AI-enabled surveillance technology may not meet the threshold for item-based export control, given its broad potential applications, but a use-based approach could restrict the export of such technology where there is reason to believe it will be used to abuse human rights.

¹⁶² Michael Brown & Pavneet Singh, *China's Technology Transfer Strategy*, Defense Innovation Unit Experimental at 29 (Jan. 2018), [https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_\(1\).pdf](https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_(1).pdf). However, there is also evidence that overall Chinese investments in the United States have been declining sharply since 2017. Rhodium Group & National Committee on U.S.-China Relations, *Two-Way Street: 2019 Update: U.S.-China Investment Trends*, http://arraysproduction-0dot22.s3.amazonaws.com/rhodiumgroup/assets/icon/RHG_TWS-2019_Executive-Summary_7May2019.pdf. A Council on Foreign Relations Task Force report notes that FIRRMA “is beginning to affect early-stage investments: some capital has moved into new sectors that are not as politically sensitive, and some dealmakers in Silicon Valley say Chinese funds are looking for deals outside the United States to avoid scrutiny. American venture capital firms are reportedly dropping their Chinese investors or walling them off, and some start-ups have forced out Chinese investors to avoid regulators.” CFR Report: Innovation & National Security at 51.

¹⁶³ Changes included expanding its coverage to include non-controlling interests (e.g., through a venture capital fund), transactions that give investors access to critical technology or sensitive data that could be exploited to harm national security, and by requiring that the review include a consideration of whether a country is of “special concern.” The Foreign Investment Risk Review Modernization Act of 2018, contained within FY2019 NDAA at 2173-2207.

¹⁶⁴ For example, the Defense Innovation Unit found that “as the U.S.-based semiconductor industry focuses on high-end designs and moves older, low-end designs offshore, the Chinese semiconductor industry now controls a significant percentage of the supply of older chips used in maintaining U.S. military aircraft and equipment designed 40 years ago and still in service.” Michael Brown & Pavneet Singh, *China's Technology Transfer Strategy*, Defense Innovation Unit Experimental at 15 (Jan. 2018).

¹⁶⁵ MINSEC is executing its first year of \$429 million in FY 2019 funding and the DoD has requested \$459 million for FY 2020. Focus areas will include specialty DoD chip needs such as radiation-hardening, secure design features, and modernizing legacy systems. DARPA also launched the Electronics Resurgence Initiative in 2017, which includes \$1.5 billion over five years for future domestic electronic systems. See *FY20 Budget Request: DoD Science and Technology*, American Institute of Physics (Mar. 28, 2019), <https://www.aip.org/fyi/2019/fy20-budget-request-dod-science-and-technology>; *A DARPA Approach to Trusted Microelectronics*, Defense Advanced Research Projects Agency, https://www.darpa.mil/attachments/Background_FINAL3.pdf; *Electronics Resurgence Initiative*, Defense Advanced Research Projects Agency (Nov. 1, 2018), <https://www.darpa.mil/news-events/2018-11-01a>.

¹⁶⁶ This government-run fund is tasked with acquiring companies along the semiconductor value chain in order to decrease China's reliance on semiconductor imports. It raised 138.7 billion RMB (about \$22 billion) in its first round in 2014. Reports from July 2019 indicate a second round has raised 200 billion RMB (about \$29 billion). However, U.S. firms control over 95 percent of Chinese market share in critical

AI-related sub-product categories (GPUs, CPUs, and FPGAs), despite significant Chinese investment in national champion semiconductor firms. See Sarah Dai, *China Completes Second Round of US\$29 Billion Big Fund Aimed at Investing in Domestic Chip Industry*, South China Morning Post (Jul. 26, 2019), <https://www.scmp.com/tech/science-research/article/3020172/china-said-complete-second-round-us29-billion-fund-will>.

¹⁶⁷ Specifically, it will be important to invest in advanced packaging formats, such as 3D stacking architectures, in order to extend the life of silicon.

¹⁶⁸ FBI Director Christopher Wray recently indicated that “nation-state actors” are “targeting academia—including professors, research scientists, and graduate students. They seek our cutting-edge research, our advanced technology, and our world-class equipment and expertise.” Remarks of The Honorable Christopher Wray, Director of the FBI, delivered at the Council on Foreign Relations, *The FBI and the National Security Landscape* (Apr. 26, 2019), <https://www.fbi.gov/news/speeches/the-fbi-and-the-national-security-threat-landscape-the-next-paradigm-shift>; see also Alex Joske, *Picking Flowers, Making Honey: The Chinese Military’s Collaboration with Foreign Universities*, Australian Strategic Policy Institute (Oct. 30, 2018); Larry Diamond & Orville Schell, *China’s Influence and America’s Interests: Report of the Working Group on Chinese Influence Activities in the United States*, Hoover Institute (2019), https://www.hoover.org/sites/default/files/research/docs/chineseinfluence_americaninterests_fullreport_web.pdf; Anastasya Lloyd-Damnjanovic, *A Preliminary Study of PRC Political Influence and Interference Activities in American Higher Education*, Woodrow Wilson Center (Sept. 6, 2018), <https://www.wilsoncenter.org/publication/preliminary-study-prc-political-influence-and-interference-activities-american-higher>; U.S. Senate Permanent Subcommittee on Investigations, *China’s Impact on the U.S. Education System* (Feb. 2019), <https://www.hsgac.senate.gov/imo/media/doc/PSI%20Report%20China's%20Impact%20on%20the%20US%20Education%20System.pdf>.

¹⁶⁹ See Letter from Kelvin K. Droegemeier, Director of OSTP, to the United States Research Community (Sept. 16, 2019), <https://www.whitehouse.gov/wp-content/uploads/2019/09/OSTP-letter-to-the-US-research-community-september-2019.pdf>.

¹⁷⁰ For example, the University of California system is auditing “risk[s] related to foreign influence” at programs funded by federal grants, and MIT has a new review process for international collaborations deemed to have “elevated risk”—including projects with connections to China. See Shailaja Neelakantan, *U of California System to Audit Campuses for Foreign Influence ‘Risk,’* Education Dive (Jul. 31, 2019), <https://www.educationdive.com/news/u-of-california-system-to-audit-campuses-for-foreign-influence-risk/559964/>; Maria T. Zuber, *New Review Process for ‘Elevated-Risk’ International Proposals*, MIT (Apr. 3, 2019), https://orgchart.mit.edu/node/27/letters_to_community/new-review-process-elevated-risk-international-proposals.

¹⁷¹ We recognize that this is easier said than done, given the opaque and complex nature of China’s research system, and the PLA’s ability to draw at will on China’s universities, companies, and overseas entities under civil-military fusion.

¹⁷² The Securing American Science and Technology Act of 2019 would “establish an interagency working group to coordinate activities to protect federally funded research and development from foreign interference, cyberattacks, theft, or espionage and to develop common definitions and best practices for Federal science agencies and grantees, while accounting for the importance of the open exchange of ideas and international talent required for scientific progress and American leadership in science and technology.” Securing American Science and Technology Act of 2019, H.R. 3038, 116th Cong. (2019),

<https://www.congress.gov/bill/116th-congress/house-bill/3038/text> (the Securing American Science and Technology Act of 2019 is currently included in the pending National Defense Authorization Act for Fiscal Year 2020).

¹⁷³ National Security Decision Directive (NSDD) 189 established that, “to the maximum extent possible, the products of fundamental research [should] remain unrestricted,” and “where the national security requires control, the mechanism for control of information generated during federally-funded fundamental research in science, technology and engineering at colleges, universities and laboratories is classification.” See NSDD 189, *National Policy on the Transfer of Scientific, Technical and Engineering Information*, National Security Council (Sept. 21, 1985), <https://fas.org/irp/offdocs/nsdd/nsdd-189.htm>.

¹⁷⁴ Such a forum could be structured similarly to the now-disbanded National Security Higher Education Advisory Board, which was established by the FBI to facilitate communication on terrorism, counterintelligence, and homeland security. See Testimony of Elsa B. Kania, Center for a New American Security, before the House Permanent Select Committee on Intelligence (July 19, 2019), <https://www.cnas.org/publications/congressional-testimony/testimony-before-the-house-permanent-select-committee-on-intelligence>; Letter from the American Council on Education et al. to FBI regarding the National Security Higher Education Board (Apr. 24, 2018), <https://www.aau.edu/sites/default/files/AAU-Files/Key-Issues/Science-Security/Letter-FBI-NSHEAB.pdf>.

¹⁷⁵ For example, in 2018 In-Q-Tel (IQT) announced new offices in London and Sydney. IQT Press Release, *IQT Establishes International Offices* (Nov. 14, 2018), <https://www.iqt.org/iqt-establishes-international-offices/>.

¹⁷⁶ The Commission would like to acknowledge the staff of the Georgetown Center for Security and Emerging Technology for submitting helpful analysis on these issues.

¹⁷⁷ As one encouraging step, in May 2019 the OECD adopted a set of principles on AI, later supported by the G20. See *Artificial Intelligence: OECD Principles on AI*, OECD, <https://www.oecd.org/going-digital/ai/principles/>. The principles support human rights and democratic values, caution against over-regulation, and offer a positive narrative about AI’s potential for social good. The Commission is also exploring the role of private and non-governmental organizations in norm development.

¹⁷⁸ Proposals such as Japan’s “Data Free Flow With Trust,” or the U.K.’s “data trusts” framework, could provide a basis, and the Commission will study these. See Remarks by Shinzo Abe, Prime Minister, delivered at the World Economic Forum (Jan. 23, 2019), <https://www.weforum.org/agenda/2019/01/abe-speech-transcript/>; Dame Wendy Hall & Jerome Pesenti, *Growing the Artificial Intelligence Industry in the U.K.* (Oct. 15, 2017), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf.

¹⁷⁹ The European Union’s AI strategy is nested within its broader emphasis on data privacy, which is most clearly manifested through the GDPR. The GDPR places restrictions on the ability of firms to collect and share personal data without consent, and provides individuals the right to revoke that consent at any given time. This privacy-first approach to data collection and sharing stands in stark contrast with the United States and Japan, which are advocating for more free flows of data. There is evidence that the GDPR has negatively impacted the competitiveness of the EU’s tech industry writ large; a recent paper from the National Bureau of Economic Research found that after the rollout of GDPR the number of venture capital deals in the tech industry within the EU declined by 26.1 percent relative to U.S.-based

firms, and the average monetary value of those deals declined by 33.8 percent. GDPR could prove to be a significant obstacle in any efforts to standardize privacy regulations, which would be a key part of any international data sharing regime. Jian Jia, Ginger Zhe Jin, and Liad Wagman, *The Short-Run Effects of GDPR on Technology Venture Investment*, National Bureau of Economic Research Working Paper No. 25248 (Nov. 2018), <https://www.nber.org/papers/w25248>.

¹⁸⁰ The Commission's view is that such a bureau would fit most naturally within the purview of the Under Secretary for Arms Control and International Security.

¹⁸¹ For example, in 2017 the United States and the U.K. signed an umbrella agreement in S&T cooperation. TS No.25/2017 Agreement between the Government of the United Kingdom of Great Britain and Northern Ireland and the Government of the United States of America on Scientific and Technological Cooperation (Nov. 30, 2017), <https://www.gov.uk/government/publications/ts-no252017-ukusa-agreement-on-scientific-and-technological-cooperation>.

¹⁸² The Five Eyes alliance traces its history to World War II, when the United States and U.K. formalized their intelligence sharing relationship under the 1946 UKUSA Agreement. Current membership includes Australia, Canada, New Zealand, the United Kingdom, and the United States. See National Security Agency, UK-USA Agreement Release 1940-1956, <https://www.nsa.gov/news-features/declassified-documents/ukusa/>.

¹⁸³ The AI Strategic Challenge will cover topics including the transition of AI technologies from research to end users, AI trustworthiness, and possible implications for international law. NSCAI has met with Five Eyes officials involved and will remain engaged with the program as it develops. See also Australian Government Department of Defence Science and Technology, *The Technical Cooperation Program*, <https://www.dst.defence.gov.au/partnership/technical-cooperation-program>.

¹⁸⁴ See Martin Dufour, *Will Artificial Intelligence Challenge NATO Interoperability?*, NATO Defence College Policy Brief No. 6 (Dec. 2018), <http://www.ndc.nato.int/news/news.php?icode=1239>; Stephan De Spiegeleire et al., *Artificial Intelligence and the Future of Defense*, Hague Center for Strategic Studies (May 17, 2017), <https://hcss.nl/sites/default/files/files/reports/Artificial%20Intelligence%20and%20the%20Future%20of%20Defense.pdf>.

¹⁸⁵ Patrick Tucker, *How NATO's Transformation Chief is Pushing the Alliance to Keep Up in AI*, Defense One (May 18, 2018), <https://www.defenseone.com/technology/2018/05/how-natos-transformation-chief-pushing-alliance-keep-ai/148301/>.

¹⁸⁶ For a similar point, see, e.g., Danzig, Technology Roulette.

¹⁸⁷ See, e.g., Paul Scharre, *Killer Apps: The Real Dangers of an AI Arms Race*, Foreign Affairs (May/June 2019).

¹⁸⁸ For a notable example of ethically-informed design, see DARPA's program on Urban Reconnaissance through Supervised Autonomy. Paulina Glass, *Here's the Key Innovation in DARPA AI Project: Ethics From the Start*, Defense One (Mar. 15, 2019), <https://www.defenseone.com/technology/2019/03/heres-key-innovation-darpas-ai-project-ethics-start/155589/>.

¹⁸⁹ For example, with respect to human rights, the United States is a party to the International Covenant on Civil and Political Rights, and with respect to the laws of war, the United States is a party to the Geneva Conventions.

¹⁹⁰ While authoritarian regimes will also seek to develop AI systems that are reliable and robust, the design and deployment of authoritarian AI systems likely will not meet the latter two criteria of ethical use and respect for rights. Authoritarian governments signing on to human rights treaties does not equate to compliance in practice. “The most important factor determining whether governments will exploit [AI surveillance technology] for repressive purposes is the quality of their governance—is there an existing pattern of human rights violations? Are there strong rule of law traditions and independent institutions of accountability?” Feldstein, AI Surveillance Paper at 10.

¹⁹¹ *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense*, Defense Innovation Board (Oct. 2019), <https://innovation.defense.gov/ai/>.

¹⁹² For example, the Department of Defense, the Department of Homeland Security, and the Intelligence Community have each articulated a set of principles to guide conduct. See Department of Defense, *Core Values*, MLDC Issue Paper #6 (Dec. 2009), <https://diversity.defense.gov/Portals/51/Documents/Resources/Commission/docs/Issue%20Papers/Paper%2006%20-%20DOD%20Core%20Values.pdf>; Department of Homeland Security, *Core Values* (July 3, 2019), <https://www.dhs.gov/core-values>; Office of the Director of National Intelligence, *Mission, Vision & Values*, <https://www.dni.gov/index.php/who-we-are/mission-vision>.

¹⁹³ This includes the perspectives of national security officials, industry leaders, academia, and civil society organizations we have consulted. The Partnership on AI, which counts over 90 institutions as members, articulates a shared commitment to fair and accountable AI. See the Partnership on AI, <https://www.partnershiponai.org/about/#our-work>.

¹⁹⁴ One example of this is the live debate on whether a moratorium is needed on facial recognition systems until legal guardrails on use and scope are developed. See, e.g., Clare Garvie, *You're in a Police Lineup, Right Now*, New York Times (Oct. 15, 2019), <https://www.nytimes.com/2019/10/15/opinion/facial-recognition-police.html>; Angelique Carson, *At House Hearing, Lawmakers want Answers on Facial Recognition from TSA, FBI*, International Association of Privacy Professionals (June 5, 2019), <https://iapp.org/news/a/at-house-hearing-lawmakers-want-answers-on-facial-recognition-from-tsa-fbi/>; Testimony of Charles H. Romine, Director of the Information Technology Laboratory, National Institute of Standards and Technology before the U.S. House of Representatives on Facial Recognition Technology (Mar. 22, 2017), <https://www.nist.gov/speech-testimony/facial-recognition-technology-frt>.

¹⁹⁵ In May and June 2019, the House Committee on Oversight and Reform held a series of hearings on the use of facial recognition technology by government and commercial entities.

¹⁹⁶ For example, rules for government access to data on Americans may need to be re-examined in light of the new kinds of insights that can be generated by AI-powered data analysis. See, e.g., Joel Brenner, *A Review of 'The Future of Foreign Intelligence: Privacy and Surveillance in a Digital Age' by Laura K. Donohue*, Journal of National Security Law & Policy at 649 (2018), http://jnslp.com/wp-content/uploads/2018/09/Review_of_The_Future_of_Foreign_Intelligence_3.pdf. As the General Counsel of the National Security Agency wrote recently: “We thought wrestling with the challenges of the Fourth Amendment in addressing electronic surveillance over the past few decades was complicated and

contentious, but setting norms for AI will surely be even more fraught with difficulty. The stakes are much higher, given that AI will be intrinsic to determinations and decisions of almost every aspect of our personal, professional and commercial lives. AI opens up the possibility of rendering intelligible for national security purposes that ocean of data.” Glenn S. Gerstell, *I Work for N.S.A. We Cannot Afford to Lose the Digital Revolution*, New York Times (Sept. 10, 2019).

¹⁹⁷ Maya Wang, *Eradicating Ideological Viruses: China’s Campaign of Repression Against Xinjiang’s Muslims*, Human Rights Watch (Sept. 9, 2018), <https://www.hrw.org/report/2018/09/09/eradicating-ideological-viruses/chinas-campaign-repression-against-xinjiangs>.

¹⁹⁸ See, e.g., Lindsay Gorman & Matt Schrader, *U.S. Firms Are Helping Build China’s Orwellian State*, Foreign Policy (Mar. 19, 2019), <https://foreignpolicy.com/2019/03/19/962492-orwell-china-socialcredit-surveillance/>; Marion Smith, *Buying Stock in These Chinese Companies Makes You Complicit in Terror on Uighurs*, Washington Post (Apr. 17, 2019), https://www.washingtonpost.com/opinions/global-opinions/think-twice-about-your-investment-portfolio-it-likely-undermines-human-rights-in-china/2019/04/17/a981b85a-6125-11e9-bfad-36a7eb36cb60_story.html; Ryan Mac et al., *US Universities and Retirees are Funding the Technology Behind China’s Surveillance State*, BuzzFeed, (May 30, 2019), <https://www.buzzfeednews.com/article/ryanmac/us-money-funding-facial-recognition-sensetime-megvii>.

¹⁹⁹ For example, in October 2019, the White House added 28 Chinese organizations to the Commerce Department entity list, including several AI companies linked to abuses in Xinjiang. *Addition of Certain Entities to the Entity List*, Bureau of Industry and Security, Dept. of Commerce, 15 FR 744 (Oct. 9, 2019), <https://federalregister.gov/d/2019-22210>. The State Department also published draft guidelines in September 2019 to aid U.S. exporters in conducting their own due diligence reviews to guard against misuse of their technology. *Draft U.S. Government Guidance for the Export of Surveillance Technology*, Dept. of State (Sept. 4, 2019), <https://www.state.gov/wp-content/uploads/2019/09/DRAFT-GUIDANCE-FOR-THE-EXPORT-OF-HARDWARE-SOFTWARE-AND-TECHNOLOGY-WITH-SURVEILLANCE-CAPABILITIES.pdf>.

²⁰⁰ Feldstein, AI Surveillance Paper at 1-2.

²⁰¹ MIT’s Lincoln Laboratory, for example, has outlined an end-to-end AI canonical architecture to identify the key enablers and possible bottlenecks in constructing the AI stack. MIT Lincoln Laboratory AI Study Report at 26.

²⁰² Milo Medin & Gilman Louie, *The 5G Ecosystem: Risks & Opportunities for DoD*, Defense Innovation Board (April 2019), https://media.defense.gov/2019/Apr/03/2002109302/-1/-1/0/DIB_5G_STUDY_04.03.19.PDF.

²⁰³ The Commission notes that updates to infrastructure are costly and may come at the expense of other technology development. In addition, 5G can consume valuable defense frequency spectrum, and may result in unevenness in global and local coverage that exacerbates inequities in access.

²⁰⁴ *The Argonne National Laboratory Supercomputer will Enable High Performance Computing and Artificial Intelligence at Exascale by 2021*, Department of Energy (March 18, 2019), <https://www.energy.gov/articles/us-department-energy-and-intel-build-first-exascale-supercomputer>.

²⁰⁵ Paper prepared by In-Q-Tel at the request of NSCAI.

²⁰⁶ Martin Giles, *Here's What Quantum Supremacy Does—and Doesn't—Mean for Computing*, MIT Technology Review (Sept. 24, 2019), <https://www.technologyreview.com/s/614423/quantum-computing-and-quantum-supremacy/>. In October 2019, Google announced that it had reached a significant milestone in demonstrating quantum viability, though the significance of the achievement has been disputed. See Frank Arute, Kunal Arya et al., *Quantum Supremacy Using Programmable Superconducting Processor*, Nature (Oct. 23, 2019), <https://www.nature.com/articles/s41586-019-1666-5>; Edwin Pednault, John Gunnels, and Jay Gambetta, *On 'Quantum Supremacy'*, IBM Research Blog (Oct. 21, 2019), <https://www.ibm.com/blogs/research/2019/10/on-quantum-supremacy/>.

²⁰⁷ Other areas of emerging technology that intersect with AI include advanced battery storage, microelectronic and semiconductors, multibeam and sensors, robotics, 3D printing, IoT hardware, and seismic imaging. Though these technologies fall largely outside the scope of the Commission's mandate, the Commission will consider the impact of developments in these adjacent fields on AI and security.

²⁰⁸ In December 2018, President Trump signed the National Quantum Initiative Act, which authorizes the government to provide \$1.2 billion over a five-year period to support a coordinated effort between federal research labs, academia, and the private sector on quantum information science. The Act directed the establishment of a National Quantum Initiative Advisory Committee (subsequently enacted by executive order in August 2019, and a National Quantum Coordination Office at the Office of Science and Technology Policy. National Quantum Initiative Act, Public Law 115-368 (Dec. 21, 2018), <https://www.congress.gov/115/plaws/publ368/PLAW-115publ368.pdf>.

²⁰⁹ In August 2016 China launched the world's first quantum satellite and has since funnelled billions of Renminbi into quantum science. Elsa B. Kania and John K. Costello, *Quantum Hegemony?: China's Ambitions and the Challenge to U.S. Innovation Leadership*, Center for a New American Security (Sept. 2018), https://s3.amazonaws.com/files.cnas.org/documents/CNASReport-Quantum-Tech_FINAL.pdf?mtime=20180912133406.

²¹⁰ Paper prepared by In-Q-Tel at the request of NSCAI. For example, in 2019 a neural network system outperformed legacy approaches in synthetic protein folding, a critical step in biological research. Paul Scharre, *Killer Apps: The Real Dangers of an AI Arms Race*, Foreign Affairs (May/June 2019).

²¹¹ Elsa B. Kania and Wilson VornDick, *Weaponizing Biotech: How China's Military Is Preparing for a 'New Domain of Warfare'*, Center for a New American Security (Aug. 14, 2019), <https://www.cnas.org/publications/commentary/weaponizing-biotech-how-chinas-military-is-preparing-for-a-new-domain-of-warfare>.

²¹² Intelligence is observed in nature beyond that of humans. Inspiring examples include phenomena such as the swarming of birds, the organization of an ant colony, and the echo location of a dolphin.

²¹³ Association for the Advancement of Artificial Intelligence, <http://aaai.org>.

²¹⁴ FY2019 NDAA at 1965.

²¹⁵ The group included John McCarthy from Dartmouth College, Marvin Minsky from Harvard University, Nathaniel Rochester with I.B.M. Corporation, and Claude Shannon with Bell Telephone Laboratories.

-
- ²¹⁶ J. McCarthy et al., *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence* (Aug. 31, 1955), <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>.
- ²¹⁷ John Launchbury, *A DARPA Perspective on Artificial Intelligence*, DARPA, slides 4-7 (Feb. 2017), <https://www.darpa.mil/attachments/AIFull.pdf>.
- ²¹⁸ *Id.*
- ²¹⁹ Y. Gil and B. Selman, *A 20-Year Community Roadmap for Artificial Intelligence Research in the US*, Computing Community Consortium and the Association for the Advancement of Artificial Intelligence (Aug. 6, 2019).
- ²²⁰ MIT Lincoln Laboratory AI Study Report at 8. The Commission added ubiquitous mobile connectivity to this list of factors.
- ²²¹ Ian Goodfellow et al., *Deep Learning*, MIT Press, Vol. 1. (2016).
- ²²² *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD: Part 3, The Deep Learning Revolution*, JASON Report, JSR-16-Task-003 at 9-25 (Jan. 2017).
- ²²³ Gary Marcus, *Deep Learning: A Critical Appraisal*, New York University (Jan. 2018), <https://arxiv.org/abs/1801.00631>.
- ²²⁴ *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD: Part 4, Deep Learning and the 'Illities'*, JASON Report, JSR-16-Task-003 at 27-32 (Jan. 2017).
- ²²⁵ See Alex Ratner, Paroma Varma, Braden Hancock, Chris Ré et al., *Weak Supervision: A New Programming Paradigm for Machine Learning*, The Stanford AI Lab Blog (Mar. 10, 2019), <http://ai.stanford.edu/blog/weak-supervision/>.
- ²²⁶ Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, MIT Press (2018).
- ²²⁷ A. Santoro et al., *One-Shot Learning with Memory-Augmented Neural Networks*, (May 19, 2016), <https://arxiv.org/pdf/1605.06065.pdf>.
- ²²⁸ Ian Goodfellow et al., *Generative Adversarial Nets*, Neural Information Processing Systems (2014), <https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- ²²⁹ Neoklis Polyzotis, Sudip Roy, Steven Euijong Whang, and Martin Zinkevich, *Data Management Challenges in Production Machine Learning*, Proceedings of the 2017 ACM International Conference on Management of Data, Chicago, IL (May 14-19, 2017).
- ²³⁰ S. Amershi et al., *Software Engineering for Machine Learning: A Case Study*, CSE-SEIP '10 Proceedings of the 41st International Conference on Software Engineering at 291-300 (2019).
- ²³¹ Dario Amodei et al., *Concrete Problems in AI Safety* (July 25, 2016), <https://arxiv.org/abs/1606.06565>.
- ²³² MIT Lincoln Laboratory AI Study Report at 27.

²³³ *Id.*

²³⁴ In-Q-Tel, <https://www.iqt.org/how-we-work/venture-capital/>.

²³⁵ *Commercial Solutions Opening*, Office of the Secretary of Defense, Defense Innovation Unit at 1 (2018), https://www.diu.mil/download/datasets/1988/DIU_CS0_-_2018_Update.pdf.

²³⁶ Geoff Orazem et al., *Why Startups Don't Bid on Government Contracts*, Boston Consulting Group (Aug. 22, 2017), <https://www.bcg.com/en-us/publications/2017/public-sector-agency-transformation-why-startups-dont-bid-government-contracts.aspx>.

²³⁷ See generally, Moshe Schwartz and Heidi M. Peters, *Department of Defense Use of Other Transaction Authority: Background, Analysis, and Issues for Congress*, Congressional Research Service (Feb. 2019).

²³⁸ *Fiscal Year 2020 Budget Request*, Comptroller & Chief Financial Officer, Office of the Secretary of Defense, Department of Defense (Mar. 2019), https://comptroller.defense.gov/Portals/45/Documents/defbudget/fy2020/fy2020_Budget_Request.pdf.

²³⁹ Table developed by Commission staff.