Final Report

National Security Commission on Artificial Intelligence

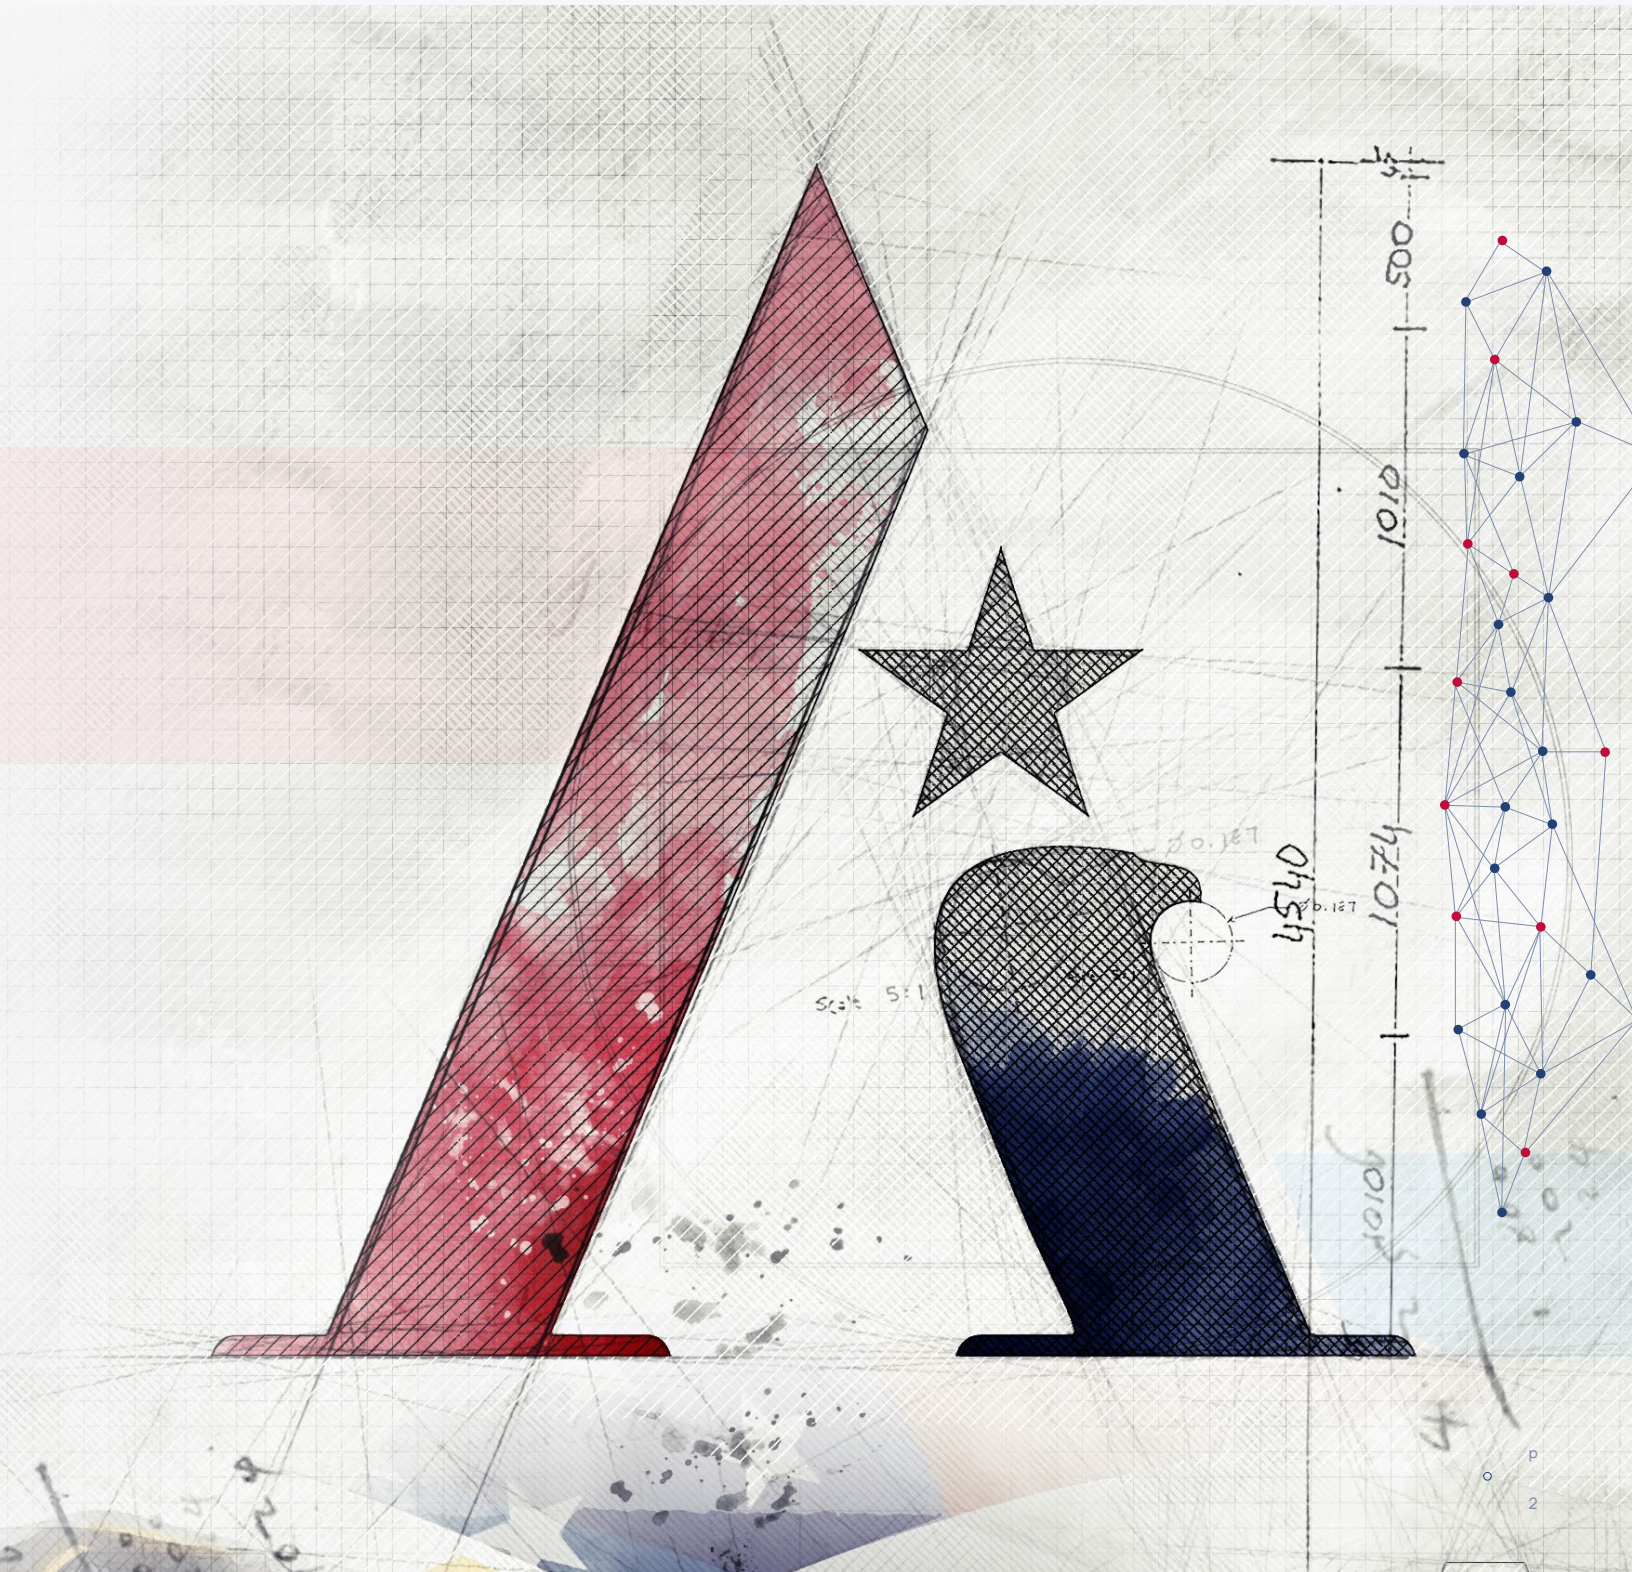

NATIONAL SECURITY COMMISSION ON ARTIFICIAL INTELLIGENCE

COMMISSION MEMBERS

Eric Schmidt
Chair

Robert Work
Vice Chair

Safra Catz

Eric Horvitz

Steve Chien

Andrew Jassy

Mignon Clyburn

Gilman Louie

Chris Darby

William Mark

Kenneth Ford

Jason Matheny

José-Marie Griffiths

Katharina McFarland

Andrew Moore

Executive Summary

No comfortable historical reference captures the impact of artificial intelligence (AI) on national security. AI is not a single technology breakthrough, like a bat-wing stealth bomber. The race for AI supremacy is not like the space race to the moon. AI is not even comparable to a general-purpose technology like electricity. However, what Thomas Edison said of electricity encapsulates the AI future: "It is a field of fields … it holds the secrets which will reorganize the life of the world." Edison's astounding assessment came from humility. All that he discovered was "very little in comparison with the possibilities that appear."

The National Security Commission on Artificial Intelligence (NSCAI) humbly acknowledges how much remains to be discovered about AI and its future applications. Nevertheless, we know enough about AI today to begin with two convictions.

First, the rapidly improving ability of computer systems to solve problems and to perform tasks that would otherwise require human intelligence—and in some instances exceed human performance—is world altering. AI technologies are the most powerful tools in generations for expanding knowledge, increasing prosperity, and enriching the human experience. AI is also the quintessential "dual-use" technology. The ability of a machine to perceive, evaluate, and act more quickly and accurately than a human represents a competitive advantage in any field—civilian or military. AI technologies will be a source of enormous power for the companies and countries that harness them.

Second, AI is expanding the window of vulnerability the United States has already entered. For the first time since World War II, America's technological predominance—the backbone of its economic and military power—is under threat. China possesses the might, talent, and ambition to surpass the United States as the world's leader in AI in the next decade if current trends do not change. Simultaneously, AI is deepening the threat posed by cyber attacks and disinformation campaigns that Russia, China, and others are using to infiltrate our society, steal our data, and interfere in our democracy. The limited uses of AI-enabled attacks to date represent the tip of the iceberg. Meanwhile, global crises exemplified by the COVID-19 pandemic and climate change highlight the need to expand our conception of national security and find innovative AI-enabled solutions.

“The NSCAI Final Report presents an integrated national strategy to reorganize the government, reorient the nation, and rally our closest allies and partners to defend and compete in the coming era of AI-accelerated competition and conflict.”

Given these convictions, the Commission concludes that the United States must act now to field AI systems and invest substantially more resources in AI innovation to protect its security, promote its prosperity, and safeguard the future of democracy. Today, the government is not organizing or investing to win the technology competition against a committed competitor, nor is it prepared to defend against AI-enabled threats and rapidly adopt AI applications for national security purposes. This is not a time for incremental toggles to federal research budgets or adding a few new positions in the Pentagon for Silicon Valley technologists. This will be expensive and require a significant change in mindset. America needs White House leadership, Cabinet-member action, and bipartisan Congressional support to win the AI era.

The NSCAI Final Report presents an integrated national strategy to reorganize the government, reorient the nation, and rally our closest allies and partners to defend and compete in the coming era of AI-accelerated competition and conflict. It is a two-pronged approach. Part I, “Defending America in the AI Era,” outlines the stakes, explains what the United States must do to defend against the spectrum of AI-related threats, and recommends how the U.S. government can responsibly use AI technologies to protect the American people and our interests. Part II, “Winning the Technology Competition,” addresses the critical elements of the AI competition and recommends actions the government must take to promote AI innovation to improve national competitiveness and protect critical U.S. advantages. The recommendations are designed as interlocking and mutually reinforcing actions that must be taken together.

Part I: Defending America in the AI Era.

AI-enhanced capabilities will be the tools of first resort in a new era of conflict as strategic competitors develop AI concepts and technologies for military and other malign uses and cheap and commercially available AI applications ranging from "deepfakes" to lethal drones become available to rogue states, terrorists, and criminals. The United States must prepare to defend against these threats by quickly and responsibly adopting AI for national security and defense purposes. Defending against AI-capable adversaries operating at machine speeds without employing AI is an invitation to disaster. Human operators will not be able to keep up with or defend against AI-enabled cyber or disinformation attacks, drone swarms, or missile attacks without the assistance of AI-enabled machines. National security professionals must have access to the world's best technology to protect themselves, perform their missions, and defend us. The Commission recommends that the government take the following actions:

Defend against emerging AI-enabled threats to America's free and open society. Digital dependence in all walks of life is transforming personal and commercial vulnerabilities into potential national security weaknesses. Adversaries are using AI systems to enhance disinformation campaigns and cyber attacks. They are harvesting data on Americans to build profiles of their beliefs, behavior, and biological makeup for tailored attempts to manipulate or coerce individuals. This gathering storm of foreign influence and interference requires organizational and policy reforms to bolster our resilience. The government needs to stand up a task force and 24/7 operations center to confront digital disinformation. It needs to better secure its own databases and prioritize data security in foreign investment screening, supply chain risk management, and national data protection legislation. The government should leverage AI-enabled cyber defenses to protect against AI-enabled cyber attacks. And biosecurity must become a top-tier priority in national security policy.

Prepare for future warfare. Our armed forces' competitive military-technical advantage could be lost within the next decade if they do not accelerate the adoption of AI across their missions. This will require marrying top-down leadership with bottom-up innovation to put operationally relevant AI applications into place. The Department of Defense (DoD) should:

First, establish the foundations for widespread integration of AI by 2025. This includes building a common digital infrastructure, developing a digitally-literate workforce, and instituting more agile acquisition, budget, and oversight processes. It also requires strategically divesting from military systems that are ill-equipped for AI-enabled warfare and instead investing in next-generation capabilities.

Second, achieve a state of military AI readiness by 2025. Pentagon leadership must act now to drive organizational reforms, design innovative warfighting concepts, establish AI and digital readiness performance goals, and define a joint warfighting network

architecture. DoD must also augment and focus its AI R&D portfolio. Readiness will also require promoting AI interoperability with allies and partners.

Manage risks associated with AI-enabled and autonomous weapons. AI will enable new levels of performance and autonomy for weapon systems. But it also raises important legal, ethical, and strategic questions surrounding the use of lethal force. Provided their use is authorized by a human commander or operator, properly designed and tested AI-enabled and autonomous weapon systems can be used in ways that are consistent with international humanitarian law. DoD's rigorous, existing weapons review and targeting procedures, including its dedicated protocols for autonomous weapon systems and commitment to strong AI ethical principles, are capable of ensuring that the United States will field safe and reliable AI-enabled and autonomous weapon systems and use them in a lawful manner. While it is neither feasible nor currently in the interests of the United States to pursue a global prohibition of AI-enabled and autonomous weapon systems, the global, unchecked use of such systems could increase risks of unintended conflict escalation and crisis instability. To reduce the risks, the United States should (1) clearly and publicly affirm existing U.S. policy that only human beings can authorize employment of nuclear weapons and seek similar commitments from Russia and China; (2) establish venues to discuss AI's impact on crisis stability with competitors; and (3) develop international standards of practice for the development, testing, and use of AI-enabled and autonomous weapon systems.

Transform national intelligence. The Intelligence Community (IC) should adopt and integrate AI-enabled capabilities across all aspects of its work, from collection to analysis. Intelligence will benefit from AI more than any other national security mission. To capitalize on AI, the Office of the Director of National Intelligence needs to empower and resource its science and technology leaders. The entire IC should leverage open-source and publicly available information in its analysis and prioritize collection of scientific and technical intelligence. For better insights, intelligence agencies will need to develop innovative approaches to human-machine teaming that use AI to augment human judgment.

Scale up digital talent in government. National security agencies need more digital experts now or they will remain unprepared to buy, build, and use AI and associated technologies. The talent deficit in DoD and the IC represents the greatest impediment to being AI-ready by 2025. The government needs new talent pipelines, including a U.S. Digital Service Academy to train current and future employees. It needs a civilian National Digital Reserve Corps to recruit people with the right skills—including industry experts, academics, and recent college graduates. And it needs a Digital Corps, modeled on the Army Medical Corps, to organize technologists already serving in government.

Establish justified confidence in AI systems. If AI systems routinely do not work as designed or are unpredictable in ways that can have significant negative consequences, then leaders will not adopt them, operators will not use them, Congress will not fund them, and the

American people will not support them. To establish justified confidence, the government should focus on ensuring that its AI systems are robust and reliable, including through research and development (R&D) investments in AI security and advancing human-AI teaming through a sustained initiative led by the national research labs. It should also enhance DoD's testing and evaluation capabilities as AI-enabled systems grow in number, scope, and complexity. Senior-level responsible AI leads should be appointed across the government to improve executive leadership and policy oversight.

Present a democratic model of AI use for national security. AI tools are critical for U.S. intelligence, homeland security, and law enforcement agencies. Public trust will hinge on justified assurance that government use of AI will respect privacy, civil liberties, and civil rights. The government must earn that trust and ensure that its use of AI tools is effective, legitimate, and lawful. This imperative calls for developing AI tools to enhance oversight and auditing, increasing public transparency about AI use, and building AI systems that advance the goals of privacy preservation and fairness. It also requires ensuring that those impacted by government actions involving AI can seek redress and have due process. The government should strengthen oversight and governance mechanisms and establish a task force to assess evolving concerns about AI and privacy, civil liberties, and civil rights.

Part II: Winning the Technology Competition.

The race to research, develop, and deploy AI and associated technologies is intensifying the technology competition that underpins a wider strategic competition. China is organized, resourced, and determined to win this contest. The United States retains advantages in critical areas, but current trends are concerning. While a competitive response is complicated by deep academic and commercial interconnections, the United States must do what it takes to retain its innovation leadership and position in the world. The U.S. government must embrace the AI competition and organize to win it by orchestrating and aligning U.S. strengths.

Organize with a White House–led strategy for technology competition. The United States must elevate AI considerations from the technical to the strategic level. Emerging technologies led by AI now underpin our economic prosperity, security, and welfare. The White House should establish a new Technology Competitiveness Council led by the Vice President to integrate security, economic, and scientific considerations; develop a comprehensive technology strategy; and oversee its implementation.

Win the global talent competition. The United States risks losing the global competition for scarce AI expertise if it does not cultivate more potential talent at home and recruit and retain more existing talent from abroad. The United States must move aggressively on both fronts. Congress should pass a National Defense Education Act II to address deficiencies across the American educational system—from K-12 and job reskilling to investing in

thousands of undergraduate- and graduate-level fellowships in fields critical to the AI future. At the same time, Congress should pursue a comprehensive immigration strategy for highly skilled immigrants to encourage more AI talent to study, work, and remain in the United States through new incentives and visa, green card, and job-portability reforms.

Accelerate AI innovation at home. The government must make major new investments in AI R&D and establish a national AI research infrastructure that democratizes access to the resources that fuel AI development across the nation. The government should: (1) double non-defense funding for AI R&D annually to reach $32 billion per year by 2026, establish a National Technology Foundation, and triple the number of National AI Research Institutes; (2) establish a National AI Research Infrastructure composed of cloud computing resources, test beds, large-scale open training data, and an open knowledge network that will broaden access to AI and support experimentation in new fields of science and engineering; and (3) strengthen commercial competitiveness by creating markets for AI and by forming a network of regional innovation clusters.

Implement comprehensive intellectual property (IP) policies and regimes. The United States must recognize IP policy as a national security priority critical for preserving America's leadership in AI and emerging technologies. This is especially important in light of China's efforts to leverage and exploit IP policies. The United States lacks the comprehensive IP policies it needs for the AI era and is hindered by legal uncertainties in current U.S. patent eligibility and patentability doctrine. The U.S. government needs a plan to reform IP policies and regimes in ways that are designed to further national security priorities.

Build a resilient domestic base for designing and fabricating microelectronics. After decades leading the microelectronics industry, the United States is now almost entirely reliant on foreign sources for production of the cutting-edge semiconductors that power all the AI algorithms critical for defense systems and everything else. Put simply: the U.S. supply chain for advanced chips is at risk without concerted government action. Rebuilding domestic chip manufacturing will be expensive, but the time to act is now. The United States should commit to a strategy to stay at least two generations ahead of China in state-of-the-art microelectronics and commit the funding and incentives to maintain multiple sources of cutting-edge microelectronics fabrication in the United States.

Protect America's technology advantages. As the margin of U.S. technological advantage narrows and foreign efforts to acquire American know-how and dual-use technologies increase, the United States must reexamine how to best protect ideas, technology, and companies without unduly hindering innovation. The United States must:

First, modernize export controls and foreign investment screening to better protect critical dual-use technologies—including by building regulatory capacity and fully implementing recent legislative reforms, implementing coordinated export controls on advanced semiconductor manufacturing equipment with allies, and expanding disclosure requirements for investors from competitor nations.

Second, protect the U.S. research enterprise as a national asset—by providing government agencies, law enforcement, and research institutions with tools and resources to conduct nuanced risk assessments and share information on specific threats and tactics, coordinating research protection efforts with allies and partners, bolstering cybersecurity support for research institutions, and strengthening visa vetting to limit problematic research collaborations.

Build a favorable international technology order. The United States must work hand-in-hand with allies and partners to promote the use of emerging technologies to strengthen democratic norms and values, coordinate policies and investments to advance global adoption of digital infrastructure and technologies, defend the integrity of international technical standards, cooperate to advance AI innovation, and share practices and resources to defend against malign uses of technology and the influence of authoritarian states in democratic societies. The United States should lead an Emerging Technology Coalition to achieve these goals and establish a Multilateral AI Research Institute to enhance the United States' position as a global research hub for emerging technology. The Department of State should be reoriented, reorganized, and resourced to lead diplomacy in emerging technologies.

Win the associated technologies competitions. Leadership in AI is necessary but not sufficient for overall U.S. technological leadership. AI sits at the center of the constellation of emerging technologies, enabling some and enabled by others. The United States must therefore develop a single, authoritative list of the technologies that will underpin national competitiveness in the 21st century and take bold action to catalyze U.S. leadership in AI, microelectronics, biotechnology, quantum computing, 5G, robotics and autonomous systems, additive manufacturing, and energy storage technology. U.S. leadership across these technologies requires investing in specific platforms that will enable transformational breakthroughs and building vibrant domestic manufacturing ecosystems in each. At the same time, the government will need to continuously identify and prioritize emerging technologies farther over the horizon.

Conclusion

This new era of competition promises to change the world we live in and how we live within it. We can either shape the change to come or be swept along by it. We now know that the uses of AI in all aspects of life will grow and the pace of innovation will continue to accelerate. We know adversaries are determined to turn AI capabilities against us. We know China is determined to surpass us in AI leadership. We know advances in AI build on themselves and confer significant first-mover advantages. Now we must act. The principles we establish, the federal investments we make, the national security applications we field, the organizations we redesign, the partnerships we forge, the coalitions we build, and the talent we cultivate will set America's strategic course. The United States should invest what it takes to maintain its innovation leadership, to responsibly use AI to defend free people and free societies, and to advance the frontiers of science for the benefit of all humanity. AI is going to reorganize the world. America must lead the charge.