November 2025

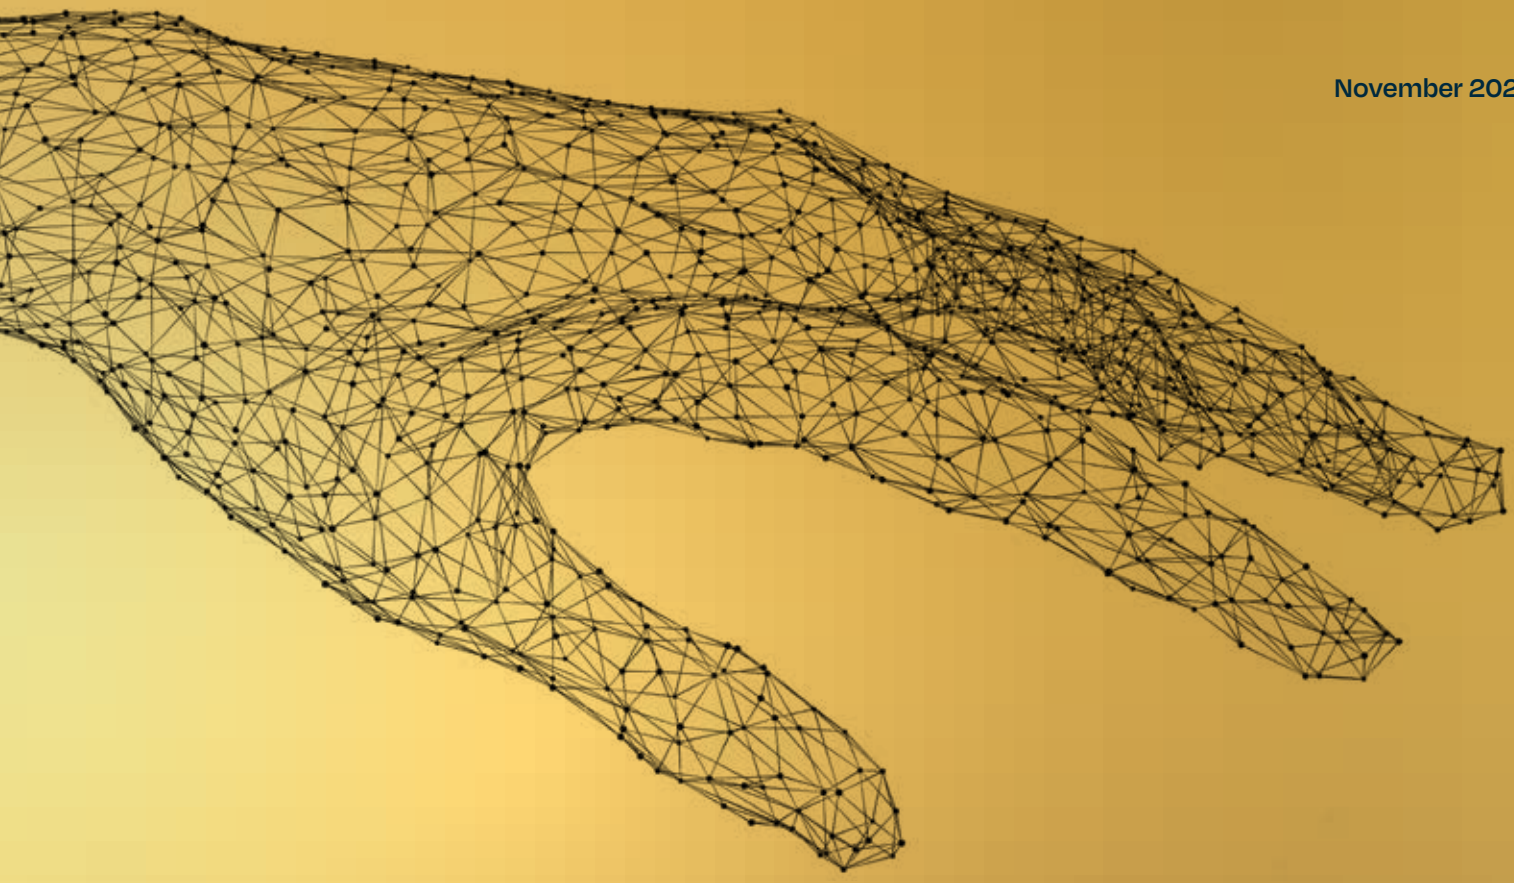

Nuclear Security Implications of AI and Emerging Technologies:

A FutureSafe Analysis of Risks and Opportunities

Douglas B. Shaw, PhD
Isabelle Williams
Patricia Jaworek
Kevin Park
Pravin Rajan

NTI
BUILDING A SAFER WORLD

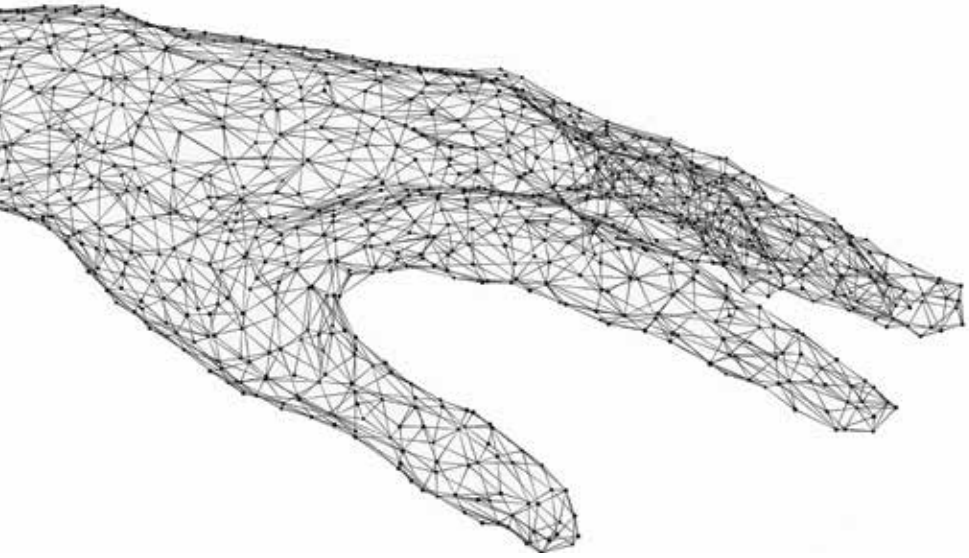

Acknowledgments

We are grateful to Nuclear Threat Initiative (NTI) President and Chief Executive Officer Christine E. Wormuth and Board of Directors Co-Chairs Sam Nunn, Ted Turner, and Ernest J. Moniz, for their vision and leadership and to former NTI President Joan Rohlfing for her important contributions to this report and the conversations and encouragement that led to it.

NTI owes a deep debt of gratitude to the experts who agreed to be interviewed for this report; their willingness to share their knowledge of complex issues and engage in productive debate makes our work possible. We are particularly grateful to NTI's Board and Science and Technology Advisory Group who engaged generatively with the work in progress multiple times and, in some cases, indulged us with multiple rounds of detailed interviews.

We are deeply grateful for generous financial support for this work from the Hess Foundation, Steven and Alison Krausz, and Ray and Meredith Rothrock.

Finally, our colleagues at NTI have made innumerable contributions to this report, all of which have been essential. We thank Carmen MacDougall, Dmitri Kusnezov, and Mark Melamed for their support and guidance; Nick Roth, Page Stoutland, James McKeon, and Lyndon Burford for their expert input; and Emma Stevens, Cathryn Crary, Annmarie Lee, Zoe Babbit, and Roberto Lachner for their crucial support. We are also indebted to NTI's Communications team, particularly Mimi Hall and Scott Nolan Smith, for their continuous guidance in shaping the report.

Douglas B. Shaw
Isabelle Williams
Patricia Jaworek
Kevin Park
Pravin Rajan

Contents

Executive Summary

This report is the first product of NTI's new FutureSafe Program on AI and Emerging Technologies, motivated by the vision for applying technology to reduce catastrophic nuclear, biological, and other advanced technological threats to human civilization.

This report builds on conversations among NTI's Science and Technology Advisory Group about how rapid commercial innovation could disrupt nuclear security in ways that might undermine a nuclear modernization approach focused on replacement of Cold War legacy nuclear arsenals. Specifically, these conversations raised concern that keeping a "human in the loop" of nuclear command and control decisions may be a necessary but insufficient approach to maintaining nuclear security against a rapidly changing background of artificial intelligence (AI) and other emerging technologies that can undermine reliability of the information on which humans make decisions.

The study team conducted interviews with 32 leading experts on nuclear security and technological innovation to identify risks and opportunities that AI and other technologies emerging from the commercial sector could pose for nuclear security. This report describes the range of responses provided in the interviews; it is not a consensus product and does not reflect the views of any single interviewee.

Interviewees provided rich and diverse responses, from which the team identified the following three risks and three opportunities:

RISKS

- Novel and newly expanded vulnerabilities of nuclear forces, particularly via attacks on the *people* responsible for nuclear security enabled by AI, big data, and networks
- New pathways to nuclear escalation
- Increased risk of nuclear proliferation and terrorism

OPPORTUNITIES

- Improved warning confidence
- Deterrence resilience and arms race stability
- New pathways to cooperative security

This study concludes that convergences of AI and other emerging commercial technologies pose novel risks and opportunities for nuclear security, and that, as a result, nuclear modernization should emphasize innovation rather than replacement of Cold War legacy architectures and systems. This report recommends that Congress (a) establish a National Security Commission on Nuclear Security Innovation and (b) create innovative vehicles for engaging cutting-edge commercial innovators in developing new capabilities and architectures that manage the risks and develop the opportunities that new technologies offer. Finally, this report can serve as a foundation for NTI to convene an expert study group on emerging technologies and nuclear security to test and build on these findings to develop a vision for national and global innovation ecosystems for preventing nuclear war into the future.

Imagine the world as mapped connections in cyberspace, creating new kinds of virtual adjacencies between objects that may be geographically separated but intimately connected digitally.

The Reason for This Study: Commercial Innovation Is Reshaping the Technological Context of Nuclear Security

Technological innovation is reshaping human civilization at astonishing clock speed. Artificial intelligence (AI) is just one example of a vast and interacting array of emerging technologies, but its rate of adoption is staggering. OpenAI's ChatGPT large language model (LLM) was adopted by more than 50 million users in less than two months, compared with 7 years for the Internet, 14 years for computers, 22 years for television, 46 years for electricity, half a century for the telephone, and 62 years for automobiles (Figure 1).[1]

FIGURE 1. **Acceleration of Technology Adoption**

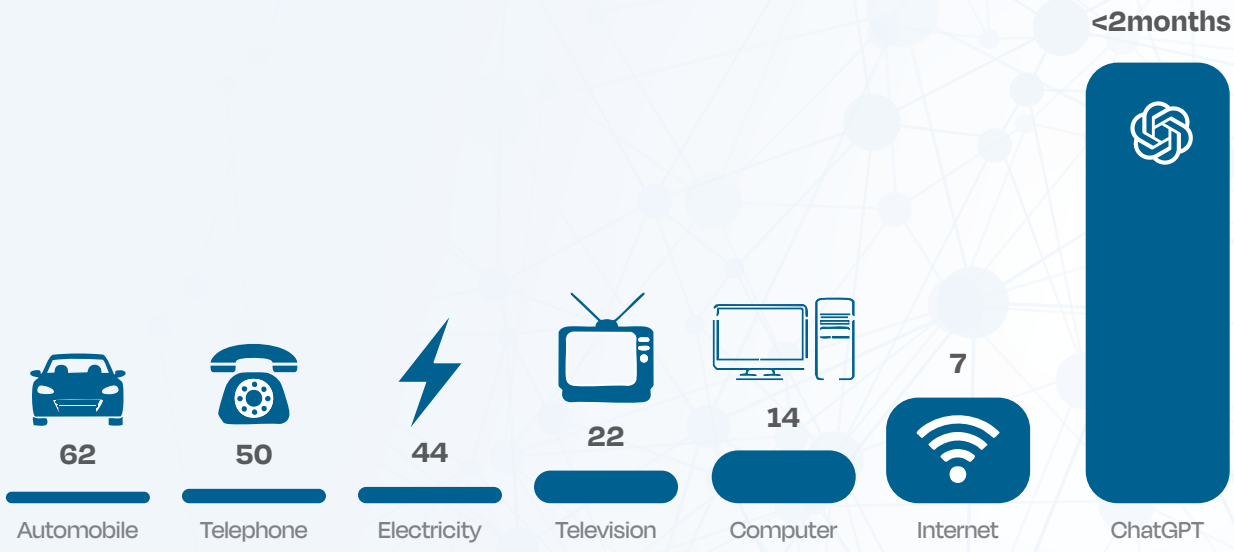

[1] Speed calculated as, 1/years to 50 million users. Data sources: Jeff Desjardins, "How Long Does It Take to Hit 50 Million Users?" Visual Capitalist, June 8, 2018 https://www.visualcapitalist.com/how-long-does-it-take-to-hit-50-million-users/#google_vignette; Krystal Hu, "ChatGPT Sets Record for Fastest-Growing User Base—Analyst Note," Reuters, February 2, 2023, https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/.

New technologies can be adopted much more quickly today because they emerge against an increasingly dynamic technological background. For example, ChatGPT is software that requires users to be familiar already with the Internet, computers, monitors, electricity, and telecommunications. As more and more technologies stack and shape each other, the range of human possibilities—including error— grows ever larger. The share of technological innovation that can stream through a wire as electrons has increased tremendously relative to the share embodied in manufactured physical objects. Marc Andreesen, a venture capitalist who invented the first web browser, famously characterized this change in 2011, observing that "software is eating the world"—the share of human activities embodied in computer code is increasing relative to the share of human activities embodied in physical objects. Software is devouring the world at an increasing rate, inevitably transforming defense along with everything else, as Trae Stephens, co-founder of defense innovation giant Anduril, observed in 2022, "software is finally eating the battlefield," and as Palantir CEO Alex Karp argues in his 2024 book *The Technological Republic*, "This next era of conflict will be won or lost with software." This report begins a challenging conversation in response to the question, *What does it mean for software to eat nuclear security?* And it explores what it means for nuclear security to be increasingly software-defined and increasingly entangled with and determined by private-sector innovation.

Nuclear security comprises the systems, tools, and practices that prevent nuclear war and control nuclear technology, including thousands of nuclear weapons on alert today that could cause catastrophic damage to human civilization in a matter of hours. This report concludes that meeting the demands of nuclear security against a backdrop of unprecedented technological dynamism demands "bleeding-edge" innovation in both tools and strategy. As Henry Kissinger, Eric Schmidt, and Daniel Huttenlocher remind us in their 2021 book *The Age of AI: And Our Human Future,* "Nuclear non-use is not an inherently permanent achievement. It is a condition that must be secured by each successive generation of leaders adjusting . . . to a technology evolving at unprecedented speed." They observe that "the management of nuclear weapons, the endeavor of half a century, remains incomplete and fragmentary" and that the "unsolved riddles of nuclear strategy must be given new attention and recognized for what they are: one of the greatest human strategic and technical and moral challenges."

The National Security Commission on Artificial Intelligence focused the needed attention of Congress and the Executive Branch on the importance of keeping a "human in the loop" of nuclear weapons command decisions. This report has a wider focus on how combinations of AI and other technologies affect the people and institutions that oversee nuclear arsenals and enact nuclear deterrence. The success of these people and institutions in preventing nuclear war hinges on accurate understandings of nuclear escalation and war termination that should be continuously updated to account for a rapidly changing technological environment. This report observes that convergences of AI and other emerging technologies accelerate known nuclear security risks; open new types of risks, such as AI and data insecurities; reduce barriers to identifying and interfering with large, geographically distributed human workforces; and create new opportunities for continuously improving nuclear security.

Scope, Method, and Examples of Underlying Technological Trends

This study examines how commercial innovations are transforming nuclear security by changing its overall technological context, rather than by making specific changes to the traditional tools of nuclear security. As such, this study is distinct in three important ways from the mainstream discourse on how AI affects nuclear security: (1) it addresses how the combination of emerging technologies (rather than just AI alone) affects nuclear security, (2) it addresses the potential adversarial application of emerging technologies to nuclear security (rather than on how AI might come to control nuclear weapons or "uplift" low-capability threat actors' efforts to acquire nuclear weapons), and (3) it foregrounds the broad trend toward greater and novel fragilities of nuclear forces and strategies in the age of AI.

The data and key findings from this report are derived from interviews with 32 senior experts in nuclear security and technology innovation (listed in the appendix) and are structured around six questions:

- What are the risks arising from the convergence of AI and other emerging technologies that increase the risk of nuclear weapons being used?
- What are the opportunities arising from the convergence of emerging technologies for preventing the use of nuclear weapons?
- In what ways might the intersection of emerging technologies make the use of nuclear weapons more or less likely?
- To what extent are key stakeholders considering the potential effects of the intersection of emerging technologies on nuclear force structure, doctrine, and policy?
- What aspects of the nuclear security implications of the intersection of emerging technologies merit more attention in the United States? Globally?
- What investments should be made (and by whom), and what governance tools should be developed (and by whom), to reduce the likelihood of and prevent the use of nuclear weapons in the age of AI?

The interviews illuminate an astonishingly dynamic technological landscape that undermines long-held assumptions at the foundation of nuclear security, including the relative invulnerability of nuclear forces to nonnuclear attack, the predictability of nuclear escalation, and the confidence that clandestine nuclear weapons acquisition programs can be found. The interviews also point toward positive opportunities, including enhanced confidence in nuclear attack warning systems, innovation-driven resilience of nuclear defense, and technologically accelerated approaches to nuclear arms control, threat reduction, and disarmament. These risks and opportunities emerge from change to the global technological environment in which nuclear security takes place, rather than conscious efforts by commercial innovators to address nuclear security concerns.

During the Cold War, a handful of governments drove nuclear security innovation with massive public investments in nuclear weapons design and manufacture, bespoke technologies such as space launch vehicles and satellites, and exquisite supercomputers that towered orders of magnitude above private-sector capabilities. Today, commercial innovation dominates many of these exquisite capabilities. For example, Elon Musk's SpaceX has more than 8,000 satellites in orbit, with plans for as many as 42,000 in one of several proposed commercial "megaconstellations." This scenario is a dramatic change from the fewer than 500 government-operated satellites in orbit during the 1990s. Reducing the cost of space launch has transformed the players and pieces of the orbital dimension of nuclear security. The changes shaped by nontraditional network-driven industries will be even greater, allowing for the collection, processing, and exploitation of unprecedented quantities of data, with implications for nuclear operations, including intelligence, surveillance, target acquisition, and reconnaissance capabilities; command of nuclear forces; and acquisition of nuclear weapons.

Humanity's accelerating pace of innovation is transforming the global technological context of nuclear security. Although nuclear weapons, delivery systems, and many other specific technologies traditionally used for nuclear security remain relatively insulated from technological change, the technological world around them has changed in ways that open new risks and opportunities for nuclear security. In addition to space commercialization, the experts interviewed for this study pointed to five converging trends driven by commercial innovation that, taken together, change the human technological environment in ways that drive revolutionary implications for nuclear security practices: (1) ubiquitous sensing and "digital exhaust," (2) large language models as "informational planets," (3) mass microtargeting, (4) swarming autonomous devices and telepresence, and (5) changes to human control. This list is exemplary, not exhaustive, but it sketches evocative dimensions of the accelerating and pervasive technological dynamism reshaping nuclear security. Each of these technological trends is described briefly in the following section to provide context for the changed technological context of nuclear security.

FIVE CONVERGING TRENDS DRIVEN BY COMMERCIAL INNOVATION

1 Ubiquitous Sensing and "Digital Exhaust"

The increasing number of sensors collecting and sharing large quantities of data almost instantly across the Internet creates new ways of identifying, locating, describing, and targeting sensitive nuclear security assets and personnel. This global sensor mesh extends from orbit into urban sewers and from our cellphones to the electric grid and wireless networks that surround us. The prevalence of accessible, portable devices such as smartphones enhances this trend by providing ready, independently maintained user interfaces for accessing and manipulating this sensor network. Microsoft AI CEO Mustafa Suleyman uses the term *digital exhaust* to refer to digital traces left by the presence and activities of people and objects. For example, when a smartphone interacts with a cell tower, it can exchange as many as 200 data fields, which can subsequently become available to the creators of apps installed on the phone or through the cellular network itself via purchase from data brokers. Ubiquitous sensing and digital exhaust combine to offer unprecedented transparency into the activities of people and organizations, including people and organizations responsible for nuclear security.

2 Large Language Models as "Informational Planets"

AI can make instant sense of this continuous data tsunami provided by ubiquitous sensors. In their 2025 book *Superagency: What Could Possibly Go Right with Our AI Future?*, Reid Hoffman and Greg Beato refer to LLMs as "informational planets" because these models contain information about everything in the world at varying degrees of specificity. Contained within each of these "informational planets," including ChatGPT, Grok, Claude, and DeepSeek, is an immense body of data. These data can reveal patterns across time, space, and quantity that elude human perception. Human behavior researchers Robert and Elaine Hubal explain that patterns of life describe "the rhythm of individuals' daily activities, and how activities are influenced by contextual factors."[2] Patterns of life apply not only to individuals but also to organizations such as nations, corporations, and nuclear weapons enterprises. The incredible speed at which machines can collect and analyze data supports the definition of more detailed signatures of behaviors and intentions. Machine learning can match these patterns nearly instantaneously at scale. LLMs can also create new pathways for data poisoning and other forms of deception. As "informational planets," LLMs can be used to observe and make inferences about anything their users choose, including the people and organizations involved in nuclear security.

[2] Robert Hubal and Elaine A. Cohen Hubal, "Simulating Patterns of Life: More Representative Time-Activity Patterns That Account for Context," *Environment International*, 172: 2023, https://www.sciencedirect.com/science/article/pii/S0160412023000260.

3 Mass Microtargeting

The information derived from commercial data can also be used to facilitate mass *microtargeting,* especially when modeling social or human organizations. Microtargeting is an advertising practice that uses data about consumers to tailor messages to each consumer. *Mass microtargeting* is the use of this technique to organize large numbers of people into target sets automatically generated by AI. Mass microtargeting can be applied to campaigns of widespread political deception as well as advertising and can employ chatbots and deepfakes. *Troll farms* (groups of online personas managed by a smaller number of users for the purpose of online deception via social media) reached an estimated 140 million Americans a month ahead of the U.S. presidential election in 2020.[3] Only the targets of these messages may be aware of their existence, meaning that adversaries can use mass microtargeting to attack large organizations by reaching the members of those organizations using personal email, text, voice calls, or social media not monitored by anyone except the targets. The "Facebook hack" was an example of a malign political application of microtargeting, during which Cambridge Analytica accessed Facebook users' activity to build psychometric profiles of users and combined those profiles to identify target groups that would be most susceptible to online manipulation. Those groups were then targeted with tailored messages to shape their behavior by boosting anger and partisan division. A novel application of mass microtargeting that crossed into the physical domain occurred in October 2024 when pagers and walkie-talkies used by Hezbollah operatives in widely dispersed locations exploded nearly simultaneously. That attack required deep penetration of a commercial supply chain to insert explosive charges into consumer products —a mixed domain attack blending cyber and physical capabilities. Mass microtargeting could be used to target people and organizations with nuclear security responsibilities for harassment, deception, extortion, or other malign purposes that could degrade their effectiveness.

4 Swarming Autonomous Devices and Telepresence

Swarming autonomous devices and telepresence allow users to deliver physical effects from great distances and at an unprecedented scale. Drone swarms can be used for search and rescue in disaster situations, autonomously coordinating among themselves to avoid obstacles, reacting to sensor inputs, and vectoring human rescuers to direct them where help is needed. Agricultural robots can deposit fertilizer or herbicides precisely on targeted plants across vast fields with minimal human supervision. The experience of turning over control of a computer to a help desk operator on the other side of an office suite or the other side of the world is increasingly common. These tools and techniques also have adversarial applications. On June 1, 2025, the Security Service of Ukraine struck five Russian airbases with drone swarms launched from civilian trucks parked near the targets, including the Belaya air base, 2,700 miles from the Ukrainian border. Nuclear-capable bombers were reportedly struck in this attack, demonstrating that swarming autonomous devices and telepresence can be used to attack important nuclear security assets.

5 Changes to Human Control

Emerging technologies can change the control relationship between human users and complex technological systems. AI tools are known to be unreliable in several ways, displaying failings such as hallucination, misalignment, and vulnerability to attack. AI hallucination occurs when a model fabricates false information that appears plausible—for example, by bolstering a law brief by citing precedents that do not exist. AI misalignment occurs when an AI tool produces results that diverge from what its human users desire or develops interests that diverge from those of its users, such as refusing to turn itself off when instructed to do so or attempting to manipulate the human users who issued the instruction to turn off. Overuse of AI tools may lead to human users becoming *deskilled* (losing their accustomed familiarity with how to perform essential but repetitive tasks that they often delegate to AI tools). As AI tools become increasingly integrated into daily life, users may be increasingly unaware of when they are using these tools. Human

[3] Karen Hao, "Troll Farms Reached 140 Million Americans a Month on Facebook Before 2020 Election, Internal Report Shows," *MIT Technology Review*, September 16, 2021, https://www.technologyreview.com/2021/09/16/1035851/facebook-troll-farms-report-us-2020-election/.

users also may think of AI tools used in their personal lives as unrelated to professional responsibilities. For example, individuals with significant nuclear security responsibilities might come to rely on agentic (a class of artificial intelligence designed to act with autonomy) AI tools to support their families, creating potential vulnerabilities if those tools stop performing as expected or are exploited for malicious microtargeting. When AI fails, it can fail in unpredictable ways, which can prevent human users from foreseeing or intervening to correct failures promptly. In a nuclear security context, such failures could be catastrophic.

Those disruptive insights point to important opportunities to improve national defense policy, force structure, and doctrine to meet future challenges. Responding to those risks and taking advantage of the opportunities would reduce nuclear risk and increase the return on investment of the current U.S. nuclear modernization program, which is expected to require between $1.2 trillion and $1.7 trillion over the coming decades.[4] This nuclear weapons modernization plan centers on replacing expensive Cold War weapons platforms—including the aging nuclear triad of bomber aircraft, intercontinental ballistic missiles (ICBMs), and submarine-launched ballistic missiles (SLBMs)—and replicating a technological architecture for preventing nuclear war designed before Moore's Law in the infancy of the software industry.

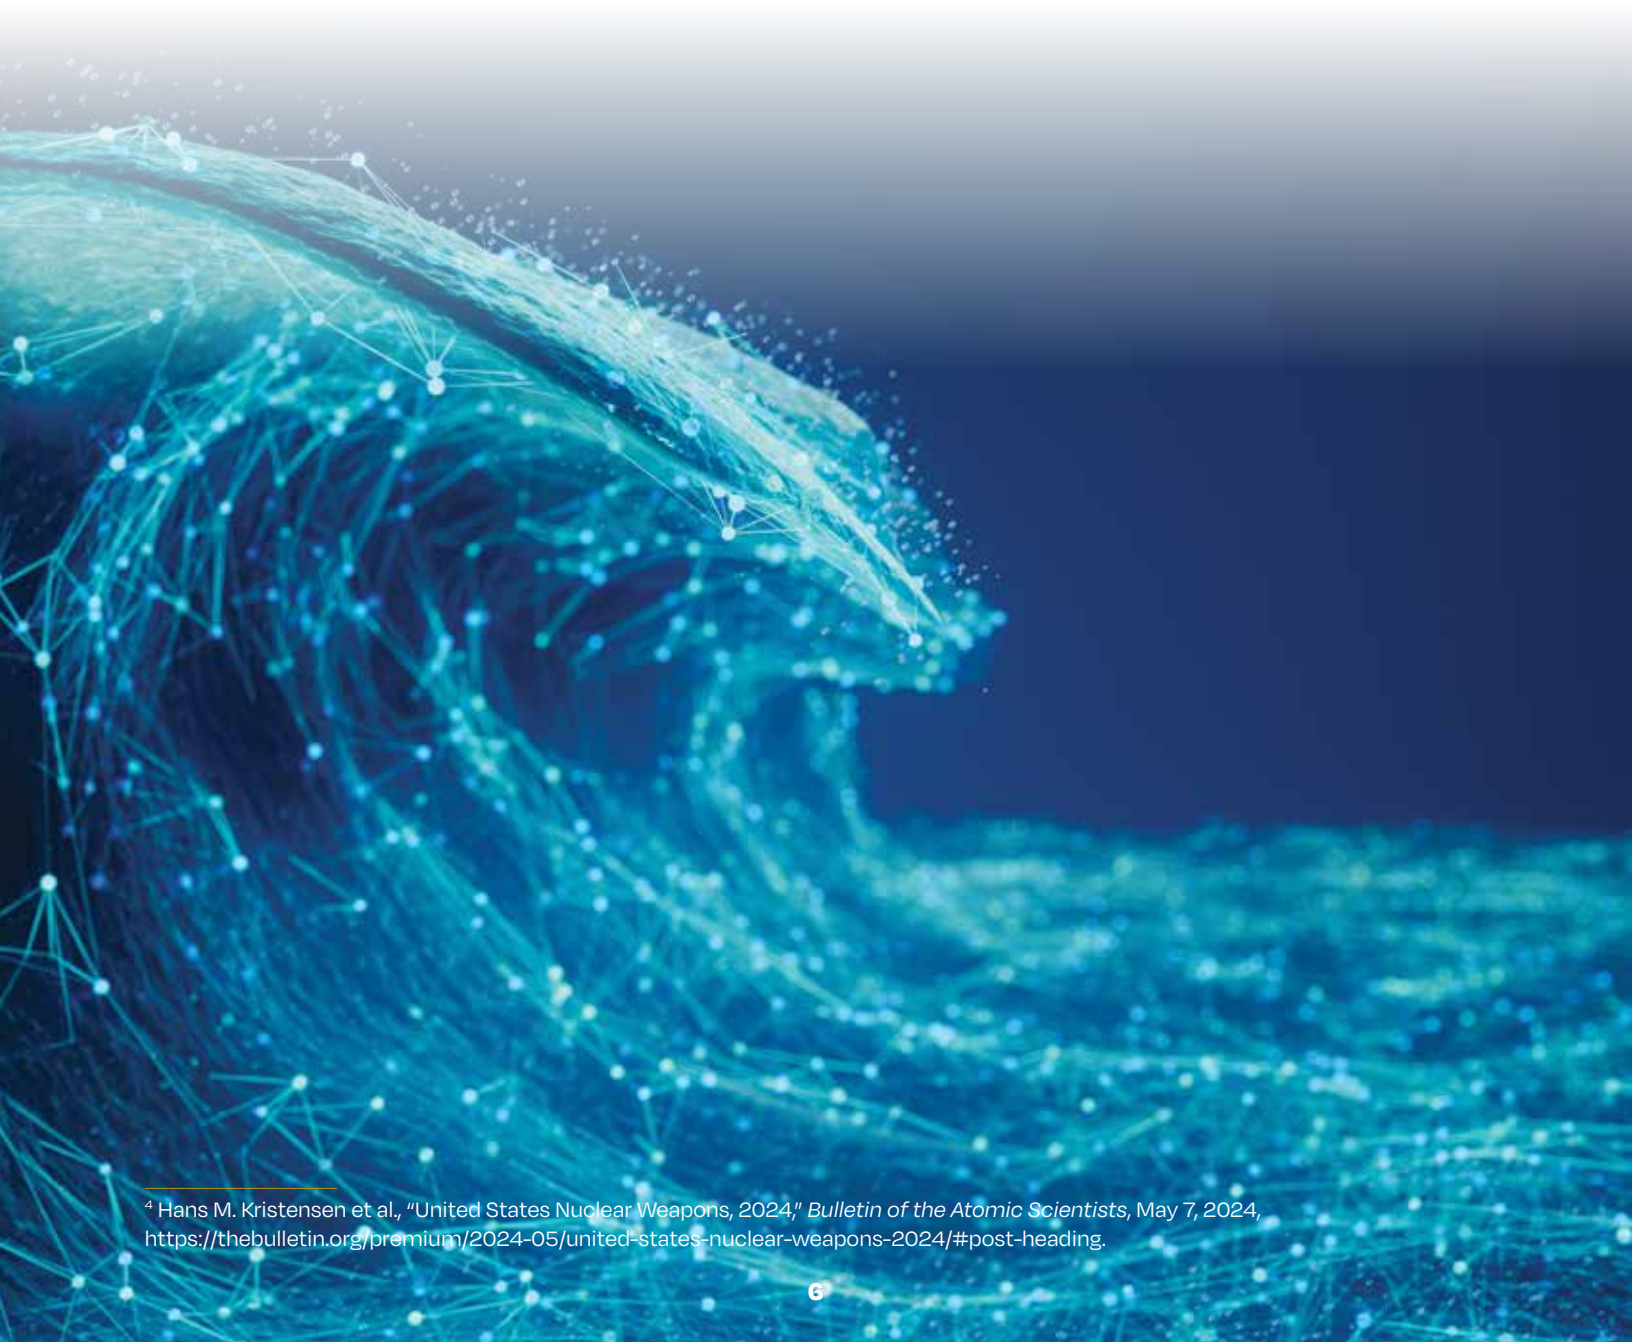

[4] Hans M. Kristensen et al., "United States Nuclear Weapons, 2024," *Bulletin of the Atomic Scientists*, May 7, 2024, https://thebulletin.org/premium/2024-05/united-states-nuclear-weapons-2024/#post-heading.

Nuclear Security **Risks** of Commercial Innovation

Interviewees identified multiple potential risks to nuclear security resulting from the changes in the global technological landscape. The following three categories reflect the most frequently raised concerns:

- Novel and newly expanded vulnerabilities to nuclear forces
- New pathways to nuclear weapons use
- Increased risk of nuclear proliferation and nuclear terrorism

Risk 1: Novel and Newly Expanded Vulnerabilities to Nuclear Forces

Numerous interviewees raised concerns that emerging technologies could make nuclear forces vulnerable in new ways or expand existing vulnerabilities. This concern is important because stable nuclear deterrence requires nuclear forces to be *survivable*—that is, sufficiently invulnerable to an adversary's first strike that leaders are confident their forces can deliver a secure second strike in retaliation (thereby deterring an adversary's first strike from ever occurring).

These new vulnerabilities are not shortcomings in the physical durability of nuclear weapons platforms, such as missiles, bombers, and submarines, or the training of human operators against known traditional threats. Rather, these new and expanding vulnerabilities are emergent properties of nuclear arsenals—globe-spanning networks of people, machines, and locations—that are increasingly embedded in the wider networks of technologies related to nuclear arsenals. For example, as LLMs evolve into "informational planets," they incidentally ingest information about the people, machines, and locations associated with nuclear arsenals.

In the past, developing an understanding of an adversary nuclear arsenal could be done only through the most resource-intensive and painstaking intelligence collection and analysis. Important analytical tools, including operations research, game theory, and scenario planning, developed partly out of these Herculean efforts early in the nuclear age. For example, the first U.S. Single Integrated Operations Plan for nuclear warfighting produced in 1960 (SIOP-62) required an analog twin of the Soviet nuclear arsenal, detailing information about targets, such as expected locations, hardness against blast, and priority. The best defense against the construction of such analog twins was secrecy.

Today, the efficacy of secrecy as partial defense against attack has changed. Exquisite, multispectral satellite imagery is available as a commercial service. Most people in the industrialized world carry cameras, microphones, and radio transmitters everywhere they go, incidentally capturing vast quantities of data. These and other innumerable streams of data flow into corporate data lakes at immense rates, becoming available to be analyzed through the lenses of LLMs and other forms of AI. Secrecy, therefore, cannot deliver the same nuclear survivability benefits it did in the past.

Several interviewees emphasized the breadth of new forms of attack, with one individual observing pointedly, "You don't have to touch a system to attack it." Another emphasized the dangers of "full-spectrum cyber" attacks, characterized by cyber-centric, multidomain operations by states capable of not only exploiting existing network vulnerabilities but also creating vulnerabilities in adversary networks. One interviewee emphasized that China and Russia have demonstrated this level of "full-spectrum cyber" capability. Another encouraged us to imagine the world as mapped connections in cyberspace, creating new kinds of virtual adjacencies between objects that may be geographically separated but intimately connected digitally.

The interviewees focused on two novel and expanding vulnerabilities of nuclear arsenals resulting from these changes to the technological environment: decision vulnerabilities and human layer targeting enabled by digital exhaust.

Decision Vulnerabilities

Nuclear deterrence depends on rationality. First and foremost, nuclear deterrence demands that national leaders make rational decisions about nuclear weapons use. In the United States, the president is the sole authority responsible for the use of nuclear weapons, making the integrity of his or her mental state and information environment central to the strict rationality on which nuclear deterrence is justified.

Malign use of social media, deceptive use of traditional media, offensive cyberattacks on support organizations and sensors, data poisoning, and targeted kinetic attacks to selectively blind or disable key capabilities are ways interviewees identified that an adversary could make coordinated, scaled attacks on the president's information environment. These capabilities could be leveraged to shape the perception and degrade the decision-making capacity of the U.S. president ahead of or during a crisis. Such a campaign could take the form of false information designed to prompt a mistaken nuclear retaliation, to undermine the president's confidence in reports of a real attack, or to undermine the president's confidence in the U.S. military or allied capabilities to retaliate against an aggressor by creating doubt about the readiness, reliability, or loyalty of components of the force. Several interviewees expressed concern that the convergence of emerging technologies could make presidential decisions and the president's information environment an attack surface, potentially leading to "good decisions based on bad information."

One interviewee observed that in the event of warning of a nuclear attack, national leaders face extreme time pressure to choose among predetermined options. Although the planning that goes into those options is sophisticated, leaders may not have deep familiarity with how they might be adjusted, channeling them toward a limited number of possible choices that might be predicted by the adversary. Another interviewee emphasized that LLM-enabled briefing materials could reinforce this channeling of presidential judgment by generating highly detailed and polished reports instantly, which could lead the president to believe that subordinates had already exercised significant human judgment about the information presented. Such materials could encourage overreliance on AI tools by a leader who has never previously faced—or perhaps even simulated—a decision to use nuclear weapons.

Malicious actors could also introduce false signatures into the information environment of national leaders responsible for nuclear-weapon-use decisions. An interviewee observed that deepfake audio or video of trusted persons, or even more complex patterns of false facts and phony advice, delivered via unsecure media could make discerning fact from fiction harder for a leader. Vast corporate data lakes supplying the massive processing capabilities of LLMs accelerate the possibility of spreading false information instantaneously across multiple channels.

Human Layer Targeting Enabled by Digital Exhaust

On New Year's Day 2023, 89 Russian soldiers were killed by a Ukrainian Armed Forces missile strike in the Russian-occupied Donetsk region. According to the Russian government, the troops died because they were using their cell phones within range of Ukrainian weapons, which "allowed the enemy to locate and determine the coordinates of the location of military personnel for a missile strike."[5] This incident illustrates how digital exhaust can inform novel modes of attack in the information age. By turning on their phones, those Russian troops doomed themselves by emitting digital signatures detailing their location and concentration.

American cryptographer Bruce Schneier said of the emerging Internet of Things (IoT), "We're building a world-sized robot." Nuclear arsenals are already world-sized robots, designed to move nuclear explosives to targets across the globe in minutes or hours. The world-spanning robots of nuclear arsenals have significant digital exhaust that could be used to target, or verify the readiness of, components of nuclear arsenals without regard to distance or traditional modes of concealment, as the story in Box 1 illustrates.

Protecting against vulnerabilities emerging from digital exhaust is not as simple as being more disciplined about keeping cell phones turned off in war zones and sensitive military facilities out of propaganda videos. The amount of digital exhaust emitted by even the most careful nuclear warfighter has increased considerably over the past two generations, and the virtual battlespace now extends far beyond the front lines. Access to live streams of data from targets or government or corporate data lakes, either through data breaches or data brokers, can enable malign actors to identify people with nuclear security responsibilities in ways that were not previously possible. Nuclear security enterprises have traditionally used personnel reliability programs and security clearance screening processes to exclude people who might have questionable judgment or loyalty, or who may be vulnerable to blackmail from such roles. This approach to force protection assumes that the vulnerability of the human layer of a nuclear arsenal is independent of adversary action. It is not robust against an environment in which computer models can monitor targets' social media behavior and learn to predict those targets' behavior better than their spouses (the highest performing human predictor in one study) by observing as few as 500 social media actions.[6] Actions similar to the "Facebook hack" or the Hezbollah pager attack could be applied to identify, target, and degrade or distract the "human layer" of a nuclear arsenal, rather than attacking the missiles, bombers, or submarines.

Personally identifiable information about these people can be used not only to identify behavior that might make individuals vulnerable to coercion but also to profile classes of targets who share characteristics or interests that could be leveraged by adversaries to manipulate, distract, or degrade the performance of large numbers of people within nuclear security contexts.

This mass microtargeting approach can inform more subtle adversarial approaches than blackmail, such as behavioral nudges—surreptitious interventions to promote desired behavior based on detailed awareness of the targets' cognitive boundaries, biases, or habits. The ability to automate the simultaneous microtargeting of large groups remotely via phishing, catfishing (when an attacker assumes a false identity to manipulate victims), or social media campaigns can allow these efforts to be both intimately tailored and sustained over time. Agentic AI also could be used to plan and conduct information operations on a large scale and to plan future kinetic or mixed-domain attacks to be executed in crisis or war. Traditional practices designed decades ago to protect against insider threats are poorly suited to defend against these novel types of information-age human layer attacks.

[5] Will Vernon and Elsa Maishman, "Makiivka: Russia Blames Missile Attack on Soldiers' Mobile Phone Use," BBC, January 4, 2023, https://www.bbc.com/news/world-europe-64159045.

[6] Wu Youyou, Michael Kosinski, and David Stillwell, "Computer-Based Personality Judgments Are More Accurate Than Those Made by Humans," *Proceedings of the National Academy of Sciences* 112, no. 4 (January 27, 2015): 1036–40, https://pubmed.ncbi.nlm.nih.gov/25583507/.

BOX 1: **DIGITAL EXHAUST IN NUCLEAR SECURITY: FROM THE SILENT GENERATION TO GEN Z**

Today's nuclear warfighters live in a different information environment from those of past generations (Launch control officer; U.S. Air Force by Airman 1st Class Brandon Valle).

In August 1969, Strategic Air Command posted its first meteorologist (the father of one of the authors of this report) to Havre Air Force Station in Montana. Born in 1944, this second-generation nuclear warfighter did not give off much digital exhaust observable to foreign spies. He was protected by two oceans and equally daunting cultural chasms separating the people of Montana's Big Sky Country from foreign adversaries. Moreover, he usually had less than $10 in his wallet, hunted for meat, exchanged sapphires he had collected from the ground for his pickup truck, and placed no more than one catalog order a year: from Sears & Roebuck at Christmas. He was difficult for Soviet or Chinese intelligence officers to observe.

Today's Gen Z and millennial nuclear warfighters live and work in a very different world. Even if they are extremely careful to limit their digital exhaust, they live in a world of cameras, microphones, and other sensors. Most have smartphones full of apps, some of which compile data about them that could be available to foreign governments either for sale by data brokers or through breaches of corporate data lakes. They are also dependent on commercial technologies in new ways, from commercial cellular networks to computerized management of complex supply chains for things such as medical care and consumer goods. Taken together, these dependencies and observable indictors of the lives of nuclear warfighters present an attack surface against nuclear arsenals very different from their Greatest and Silent Generation forebearers, around whose protection the system was designed.

Risk 2: New Pathways to Nuclear Weapons Use

Several interviewees mentioned the potential of new pathways to nuclear weapons use as a major risk of AI and other emerging technologies for nuclear security.

Outer Space Commercialization

Emerging technologies such as reusable boosters and proliferated low Earth orbit architecture have dramatically decreased the cost of space launch and satellite communications, making these once-scarce assets more affordable (Figure 2). This circumstance has transformed space activities by increasing society's reliance on orbital architecture; widening the potential for counterspace, sabotage, and ground attack; and weakening the defenses of legacy military space systems.

FIGURE 2. **Commercial Innovation Is Making Space Launch Less Expensive**

More Boost for the Buck: Increase in mass (in kg) $1 million can launch to Low Earth Orbit (in constant FY21 dollars)

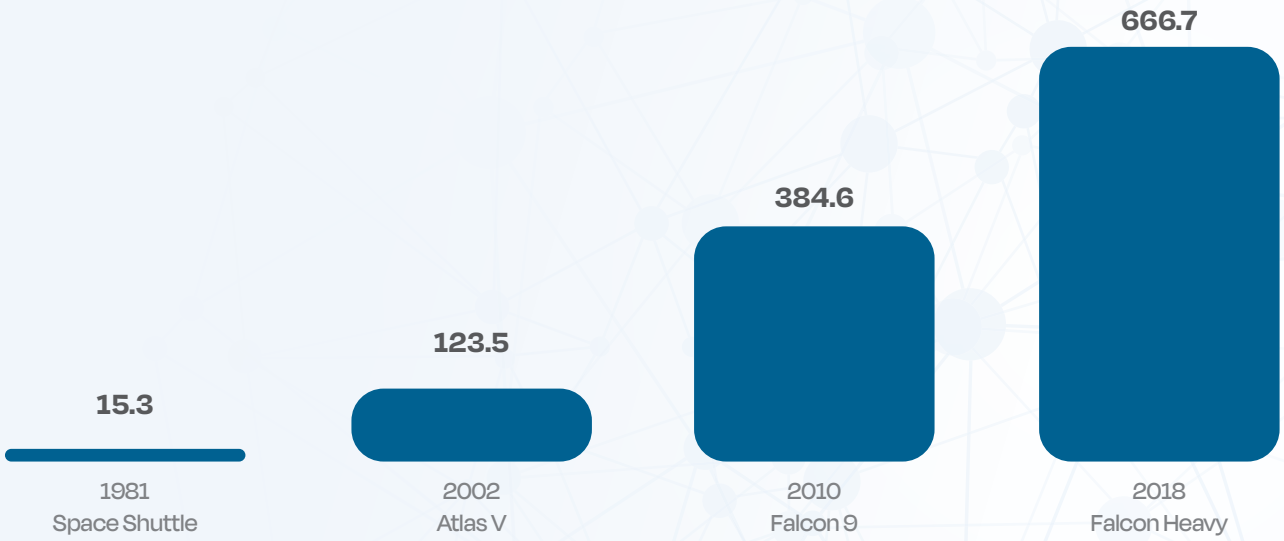

Data Source: Thomas G. Roberts, "Space Launch to Low Earth Orbit: How Much Does It Cost?," Center for Strategic and International Studies Aerospace Security Project, September 1, 2022, https://aerospace.csis.org/data/space-launch-to-low-earth-orbit-how-much-does-it-cost/.

Today's most advanced militaries rely heavily on space-based assets. U.S. warning and communications capabilities have relied for decades on large, expensive satellites in high orbit. Commercial mega constellations, such as Starlink, now provide consumers dual-use space capabilities that were once limited to major powers. The U.S. government increasingly uses SpaceX's defense product Starshield. Those consumer capabilities enable more effective battle networks, combining sensors, command and control, and kinetic effects into coordinated grids of observation, orientation, decision-making, and action. Commercial space assets, however, could also interfere with military capabilities by colliding with military satellites or creating debris that subsequently collides or interferes with military satellites.

The development of large commercial satellite networks has also transformed the benefits that can be derived from space. For example, satellite phone service that once cost thousands of dollars has been transformed into a service that can be accessed from today's ubiquitous smartphones.

This public-private entanglement has emerged in new ways in Russia's war with Ukraine. Dmitri Alperovitch, co-founder and chairman of the Silverado Policy Accelerator, observed that SpaceX's Starlink service has been "essential" to Ukraine's war effort.[7] SpaceX CEO Elon Musk has opined publicly on the possibility of Russian nuclear escalation in that conflict, tweeting "If Russia is faced with the choice of losing Crimea or using battlefield nukes, they will choose the latter."[8] And Russia may be developing new counterspace capabilities, reportedly including the Cosmos 2553 satellite, which some have speculated could be a test-bed for a nuclear anti-satellite weapon."[9]

The preceding example points to a deep technological imbalance shaping escalation in space. In the words of one interviewee, Russia "has lots of military hardware, great space launch capabilities, good sensor capabilities, and has been successful getting microchips and other parts despite sanctions." However, it lacks the people needed to compete in emerging technologies over the long run. As a result, Russia may plan to use nuclear weapons as a low-technology equalizer against exquisite U.S. capabilities. Placing a nuclear weapon in space would violate a key provision of the 1967 Outer Space Treaty, potentially undermining international confidence in that agreement and opening the door to wider militarization of outer space.

In May 2025, U.S. President Donald Trump announced plans for a "Golden Dome missile defense shield" system for the stated purpose of "ending the missile threat to the U.S. homeland." He said that it would achieve a "success rate very close to 100 percent" by deploying "next generation technologies across the land, sea, and space, including space-based sensors and interceptors."[10] That claim is extraordinary, and the capabilities of the proposed Golden Dome system have yet to be demonstrated. Moreover, Russian, Chinese, and North Korean planners who threaten the United States with nuclear missiles today may plan for a future in which those missiles will be less effective by looking for other ways to threaten the United States, making the future net implications of a Golden Dome uncertain. During the Golden Dome announcement, President Trump explained that "Ronald Reagan wanted it many years ago, but they didn't have the technology." Private companies developed much of the relevant new technology available today that was unavailable during the Reagan administration, including reusable boosters, which reduce the cost of launching payloads to orbit, and proliferated low Earth orbit architectures such as the one demonstrated by Starlink. In the 1980s, maintaining thousands of satellites in orbit seemed too difficult and expensive for the U.S. government; today, SpaceX is delivering high-speed, low-latency satellite Internet to consumers globally from a commercial megaconstellation.

[7] Shania Shelton, "NASA Chief Calls for Investigation into Report that Musk and Putin Have Spoken Regularly," *CNN Politics*, October 25, 2024, https://www.cnn.com/2024/10/25/politics/elon-musk-vladimir-putin.

[8] Tristan Bove, "Elon Musk Supports Russia Keeping Crimea—Because He's Worried About Nuclear Escalation and World War III," *Fortune*, October 17, 2022, https://fortune.com/2022/10/17/elon-musk-world-war-3-could-happen-russia-nuclear-response-crimea-putin/.

[9] Theresa Hitchens, Is Russia's Cosmos 2553 Satellite a Test for a Future Orbital Nuclear Weapon?" *Breaking Defense*. May 22, 2024, https://breakingdefense.com/2024/05/is-russias-cosmos-2553-satellite-a-test-for-a-future-orbital-nuclear-weapon/.

[10] The White House, "President Trump Makes an Announcement with the Secretary of Defense," May 20, 2025, https://www.whitehouse.gov/videos/president-trump-makes-an-announcement-with-the-secretary-of-defense/.

Erosion of the Nuclear Threshold

Nuclear weapons have not been used in combat for more than 80 years. Some theorize that this norm of non-use results from a taboo; others disagree. Setting aside the motivation for restraint from nuclear weapons use, maintaining large, geographically distributed nuclear arsenals that are *always* ready for use but *never* activated accidentally is a monumental technological accomplishment. The continuous achievement of *always/never* control over nuclear weapons by multiple governments is a permanent necessity for national and global security in a world that includes nuclear weapons.

During the Cold War, deterrence theorists sought "first-strike stability"—the strategic condition in which neither side would gain a decisive advantage by using nuclear weapons first. Numerically large, geographically distributed, diverse, and hardened nuclear forces promote first-strike stability by increasing the physical difficulty and planning complexity for adversaries that might wish to attempt to disarm a nuclear foe. In this context, some weapons systems—such as ICBMs, with multiple, independently targetable nuclear warheads—were considered more destabilizing than others because if they were struck in a surprise attack on the ground, 1 or 2 attacking nuclear weapons might destroy 10 or more retaliatory weapons.

Technological innovation may be making strategic nuclear delivery vehicles and nuclear command and control systems more vulnerable to nonnuclear attack, undermining first-strike stability. Interviewees in the study identified a number of ways that rapid or instantaneous novel nonnuclear strike systems could be used to attack nuclear weapons or command and control systems. Millimeter-wave radar and other advanced targeting systems can improve the accuracy of missile systems, potentially undermining the assumed invulnerability of hardened targets to conventional attack. Advanced cyber- and mixed-domain attacks can harm, distract, or confuse the personnel constituting a nuclear arsenal without regard to their number or geographic distribution; one interviewee described the potential of "digital decimation" attacks planned over years to degrade retaliatory capabilities by 10 percent or more when activated. In the future, weaponized civilian infrastructure, such as the 2024 Hezbollah pager attack, could be used to degrade nuclear retaliatory forces. Hypersonic missiles, suborbital bombardment, or drone attacks could change the warning signatures of attacks on nuclear forces. After the interviews for this study concluded, a novel vulnerability of nuclear forces to conventional attack was demonstrated in Russia when Ukraine attacked five Russian air bases with drones launched from trucks carried near their targets on public roads. Nuclear-capable aircraft were reportedly destroyed or damaged in those attacks.

Rapid or instantaneous strike systems can undermine warnings of conventional and nuclear attacks by enabling conventional means of striking nuclear targets. That occurrence confuses a situation in which intercontinental-range missile systems have largely been segregated to avoid confusion about whether they are carrying nuclear or conventional payloads and opens the door to runaway escalation from a misinterpretation of an attack warning. That misinterpretation increases the prospect of a strategic first strike that would leave a defender with a diminished nuclear retaliatory capability. As a result, the separation between conventional and nuclear warfare becomes increasingly blurred.

Risk 3: Increased Risk of Nuclear Proliferation and Nuclear Terrorism

Interviewees shared their concerns that the convergence of emerging technologies could lower technological barriers to nuclear proliferation and nuclear terrorism. Multiple interviewees expressed serious concerns about the combination of AI and advanced manufacturing techniques easing access to weapons-usable nuclear materials, especially through clandestine uranium enrichment. However, they also emphasized the difficulties in creating and maintaining a nuclear weapons enterprise.

Access to Uranium Enrichment

The most significant technological barrier to building a nuclear weapon is the acquisition of sufficient quantities of weapons-usable nuclear materials—notably, plutonium or highly enriched uranium (HEU). This barrier is large because although natural uranium is abundant in the Earth's crust, only about 0.7 percent comprises the fissile isotope U–235, which must be refined and separated from the more abundant U-238 to be usable in a nuclear weapon.

Although the barriers to enriching uranium remain high, enrichment technology has improved, lowering financial and energy requirements. Those technological advancements have allowed multiple countries to enrich uranium on a commercial basis under International Atomic Energy Agency (IAEA) safeguards. Those improvements also reduce the signatures of enrichment, however, making it easier to hide clandestine activities. One interviewee stressed that the enrichment requirements of a civilian nuclear power program are vastly greater than those needed for a clandestine nuclear weapons program: fueling a single 1,000-megawatt civilian light water reactor for a year typically requires about 100,000 or more separative work units (SWU—a measurement of uranium enrichment), many times the rough-order-of-magnitude of the enrichment required one time for a single nuclear explosive. Owing to the lower quantities necessary, a smaller enrichment capability than what is considered industrial scale could be sufficient to enable an improvised nuclear explosive device or even a small nuclear arsenal.

Emerging technologies could make uranium enrichment more accessible by enabling the use of autonomous discovery to develop sub-industrial scale, low-observable pathways to these technologies. The processes of enrichment are derived from scientific observation and include several known technical pathways. Multiple interviewees expressed concern that AI could help malign actors access those technical pathways, discover additional pathways to uranium enrichment, and facilitate their clandestine use. AI could be used to make esoteric and tacit knowledge more accessible or to speed independent scientific inquiry by identifying dead ends and helping to design research programs. Lowering the technical barrier to uranium enrichment could not only lead to the spread of nuclear weapons but also undermine the progress made in recent decades toward nuclear forensics for post-detonation attribution of the source of fissile material used in a nuclear explosive.

Nuclear Security **Opportunities** of Commercial Innovation

Interviewees most frequently cited three areas of opportunities for nuclear security driven by the convergence of emerging technologies:

- Improved warning confidence
- Deterrence resilience and arms race stability
- Nuclear arms control, nonproliferation, and threat reduction

Opportunity 1: Improved Warning Confidence

Interviewees frequently pointed out the potential for emerging technologies to enhance the reliability of warnings of attacks on nuclear forces, including with more sensors and advanced data fusion. By reducing uncertainty during crises, minimizing false warnings of nuclear attack, and increasing opportunities for de-escalation, those technologies could help lower the risk of nuclear catastrophe. They also could create a greater sense of security and strengthen strategic stability by increasing the accuracy and timeliness of warnings while reducing false warnings. AI-enabled warning systems could enhance current capabilities by leveraging new sensor technologies, quantum systems, and unmanned platforms. AI also can rapidly interpret vast amounts of data from early-warning systems more effectively than humans, recognizing patterns and anomalies.

Rapid analysis of more data from diverse sources should enable not only more precise tracking of adversary actions but also deeper insights into their underlying intent. This knowledge is crucial in the context of potential warnings of nuclear attack, when decision-makers could have only minutes to determine whether to use retaliatory nuclear forces. For example, the U.S. ICBM force is "deployed in hundreds of nuclear-hardened silos and can be launched to reach targets within minutes, creating a complex targeting problem for adversaries."[11]

Given that a large-scale use of nuclear weapons has never occurred, synthetic data would be needed to achieve the projected benefits in warning confidence (artificial data created by algorithms to simulate events that have not happened). One interviewee stressed the need for great care in the development, validation, and application of synthetic data to a task as consequential as nuclear attack warning. Identifying appropriate standards and safeguards for such synthetic data would be an important first step. Those standards and safeguards could support greater confidence in political leaders and the public and could potentially be used as a basis for risk-reduction negotiations with potential nuclear adversaries, and to demonstrate restraint to non–nuclear-weapon-state parties to the Treaty on the Non-Proliferation of Nuclear Weapons (NPT).

[11] Office of the Deputy Assistant Secretary of Defense for Nuclear Matters, "Chapter 3—Nuclear Delivery Systems," *Nuclear Matters Handbook*, 2020 (revised), https://www.acq.osd.mil/ncbdp/nm/NMHB2020rev/chapters/chapter3.html.

Biases and other possible shortcomings of AI models could be overcome by implementing voting logic among multiple, distinct models. A "multi-model rule" for early-warning AI implementation can mitigate the danger of AI hallucinations, biases, or other errors that can result in false warnings by combining the output of multiple models.

AI also could facilitate understanding of the human layer of nuclear arsenals. Analysis of whom leaders are talking to and who exercises influence within the national command authority of nuclear arsenals can help illuminate rising dangers. One interviewee observed that "the [COVID-19] pandemic was extremely problematic" in terms of understanding Russian nuclear decision-making "because Putin is a germaphobe and didn't talk to anyone."

Multiple interviewees also stressed that transparency creates an advantage for open societies in the long run. Closed societies often depend on a greater volume of secrecy, not just about military plans and technology but also about grievances of elements of civil society and the methods by which those grievances are repressed. In such cases, facts about a state's tools of coercion used on its own citizens can become weapons against repressive regimes.

Opportunity 2: Deterrence Resilience and Arms Race Stability

The survivability of nuclear retaliatory forces—referred to as a secure second-strike capability—is a key variable in nuclear deterrence stability. In the 1950s, pioneering deterrence theorist Albert Wohlstetter transformed U.S. nuclear deterrence policy by asking how vulnerable U.S. nuclear forces (rather than U.S. cities) were to a Soviet first strike. That question led to a series of actions intended to make U.S. nuclear forces survivable against a potential future Soviet attempt at a disarming first strike. Whether decision-makers focus sufficient attention today on how technological change is either creating or reducing vulnerabilities of nuclear forces to potential first strikes is unclear. For example, today, leaders recognize that military aircraft parked "wingtip to wingtip" are more vulnerable to bomb, missile, or drone attack than aircraft that are geographically dispersed. Because information accounts for an increasing share of nuclear deterrence, could vulnerabilities for nuclear retaliatory capabilities emerge from the configuration of data in cyberspace being stored wingtip to wingtip in a virtual sense?

Deterrence resilience is conceptually much wider than the survivability of specific nuclear weapon platforms. Overall, it is created by strengthening strategic forces and developing ways to better respond to adversarial attempts to attack, degrade, or disrupt deterrent capability.

Interviewees described paths to improve nuclear command, control, and communications (NC3) resilience against attacks; to expand nuclear force protection to include data; and to leverage mixed-reality, blockchain, and advanced manufacturing techniques to lower the costs and risks of maintaining a nuclear arsenal.

Multiple interviewees also suggested creating new mechanisms for engagement between the nuclear security enterprise and industries responsible for supporting strategic forces mission assurance. One interviewee suggested that, in some circumstances, national security may benefit from a government request that certain private-sector actors change their data-management practices, either generally or in specific response to a crisis or incident.

Deterrence Resilience

Nuclear deterrence comprises three elements: the military **capability** to punish or deny any benefit of offensive action to an adversary; the political **credibility** to execute that threat; and the **communication** of that threat to the adversary. As the number and types of attack surfaces expand, adversaries are developing new capabilities and doctrines designed to destroy, degrade, or disrupt one or more of those elements. The target of these capabilities is not just the triad of delivery systems but the associated data, assets in outer space, and political cohesion needed to conduct the practice of deterrence.

The first step toward deterrence resilience is to fortify NC3 against all forms of potential adversarial attack, including those driven by technological developments. For example, the United States could transition from vulnerable, legacy military space architectures to approaches that take advantage of proliferated low Earth orbit architectures, such as Starlink. Such constellations not only are more resilient but can be further bolstered with "responsive launch" and in-space mobility capabilities that enable rapid replacement of destroyed systems.

New technological and doctrinal measures can also be leveraged to better protect the human layer of nuclear deterrence (Box 2). In interviews, President Emeritus of Rensselaer Polytechnic Institute Shirley Ann Jackson, described the potential to create a "digital immune system" that could protect decision-makers, nuclear warfighters, and employees of the nuclear enterprise from data-driven attacks. She envisioned the creation of "sentinels of information age attack," taking the form of "digital avatars"—decoys of members of a target class, such as a missile launch control officer, that could be inserted into data lakes to monitor and provide warning of attacks on that target class.

BOX 2: **DIGITAL RESILIENCE**

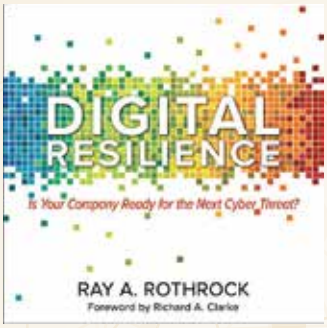

Nuclear arsenals are vast networks of people and machines, making them vulnerable to network attacks. Emerging technologies can help bolster deterrence resilience in nuclear security by enabling novel ways of verifying deceptive information and of shielding against human layer attacks, including through increased digital resilience. In his 2018 book *Digital Resilience*, founder and CEO of FiftySix Investments Ray Rothrock observes that "The risk that your network will be attacked depends, first and foremost, on its size (number of connected users, or 'nodes'). The bigger the network, the greater the risk. . . . Moreover, because today's Internet encompasses a trillion-plus-node Internet of Things, the line separating digital resilience from the more familiar physical resilience of organisms, people, corporations, commercial aircraft, governments, and ecosystems is rapidly dissolving." Achieving digital resilience is not solely an information technology problem achieved by simple software or hardware solutions; it requires a systemwide organizational strategy to establish multidomain, full-spectrum cyber defenses.

Countering Deception

As deception becomes cheaper and more widespread, adaptive AI defenses may be the only path to counter AI deception. Knowing with certainty about information sources—such as data used to train AI—will become more important to build deterrence resilience. New information management systems such as blockchain, a type of distributed ledger technology, can be integrated with sensors and big data to help combat AI-accelerated deception by enabling cryptographic verification of exactly what data were used to train AI models.

RAND researcher Edward Geist foresees that the rise of AI will transform the offense-dominant environment shaped by nuclear weapons (in which a dedicated attack from a large nuclear arsenal cannot be stopped) into a deception-dominant environment, in which the introduction of false data into adversary systems will become easier than the ability to detect said false data. Such a dramatic change to the technological substrate of nuclear deterrence would significantly undermine the coherence of and confidence in nuclear deterrence as practiced today. In his book *Deterrence Under Uncertainty: Artificial Intelligence and Nuclear Warfare*, Geist observes that Russia is already adapting its defense acquisitions to exploit a deception-dominant environment.[12]

Nuclear war planning has long depended on building elaborate models of adversaries' military capabilities, assets, preferences, and likely behaviors. During the Cold War, nuclear-armed states planned hypothetical nuclear wars on the basis of what today would be considered very sparse information, such as rough estimates of the numbers and locations of adversary nuclear weapons systems, which can be described as *analog twins*. Today, the growing use of highly detailed *digital twins* (virtual replicas of complex systems) enables the synthesis of larger volumes of data about adversary capabilities and intentions. Such synthesis includes creating digital twins of not only nuclear arsenals but also the environmental, supply chain, and social factors that support them.

Emerging technologies are also providing potential alternatives to nuclear weapons. Nonnuclear weapons, both kinetic and nonkinetic, could fulfill many of the missions traditionally reserved for nuclear strikes, reducing—or, in some cases, eliminating—the unique military utility of nuclear weapons. Fiona Cunningham, in her book *Under the Nuclear Shadow: China's Information Age Weapons in International Security*, observes that China is seeking deterrence leverage against the United States by threatening U.S. conventional capabilities with precision missiles, offensive cyber, and counterspace capabilities to constrain U.S. military options.[13] Understanding attempts by potential nuclear adversaries to manipulate U.S. nuclear weapon decisions is essential to preventing nuclear war and could form the foundation for a wider approach to negotiated nuclear risk reduction.

Crisis communication channels with adversaries and mechanisms for war termination should be modernized with the same priority and investment as are applied to nuclear weapon platforms. The same technologies used to assess adversary intent and predict behavior also can be used to improve communication during a crisis, helping to identify mutually acceptable solutions that might not occur to a human. LLMs can rapidly model potential adversary reactions, increasing confidence in the resilience of policy options across a wide range of scenarios. Whereas the use of nuclear weapons might stem from the failure of human imagination to consider different possibilities, the capacity of AI to rapidly explore a wide range of options for conflict resolution and diplomacy could expand the set of viable alternatives and reduce the risk of escalation.

As private-sector entities play an increasing role in shaping nuclear security, new forms of public-private partnership will be important. Those forms could include steps such as increased outreach by governments to corporate stakeholders or reshaping of corporate entities in ways that privilege strategic stability and war prevention. Multiple interviewees recommended including relevant private-sector actors, such as AI labs and companies with large data lakes, in frequent wargaming with potential conflict scenarios. This pursuit would benefit not only the private-sector participants by familiarizing them with the possible consequences of conflict but also the official participants by familiarizing them with the capabilities that the private sector can apply to complex technological problems.

[12] Edward Geist, *Deterrence Under Uncertainty: Artificial Intelligence and Nuclear Warfare* (Oxford University Press, 2023).

[13] Fiona Cunningham, *Under the Nuclear Shadow: China's Information Age Weapons in International Security* (Princeton University Press, 2025).

Arms Race Stability

Interviewees surfaced ways that emerging technologies could be applied to dampen incentives for arms racing and to control the significant costs associated with maintaining a nuclear arsenal. One interviewee emphasized that the central challenge of arms race stability is to convince adversaries of "the things we are not going to do." If nuclear adversaries believe that the other is not capable of restraint, *they* will not consider restraint.

Another interviewee envisioned a future U.S. nuclear weapons complex that is fully digitized. This approach builds on the success of capabilities such as the National Ignition Facility at Lawrence Livermore National Laboratory (LLNL) that has enabled the U.S. government to gain insights and make decisions about nuclear weapon programs that would not have been possible with nuclear explosive testing. LLNL has pioneered other simulation approaches to improve confidence in the nuclear stockpile—for example, by creating a Polymer Production Enclave that mirrors some of the tools at the Kansas City National Security Campus, allowing LLNL personnel to better understand how they function. "We tend to think of production as a one-way arrow," reflected an interviewee, "but thinking about it as a system . . . and adding digital twins can support optimization tools that can lead to new approaches."

In 2006, the National Nuclear Security Administration offered a vision that included elements of that approach in *Complex 2030: An Infrastructure Planning Scenario for a Nuclear Weapons Complex Able to Meet the Threats of the 21st Century.* That vision prioritized a "responsive infrastructure" to produce tailored nuclear weapons rather than maintaining larger stockpiles of legacy weapons.

One interviewee observed that innovation today has opened the door to making "things on a timescale and cost that shakes our adversaries' calculus" and becoming "fleet enough of foot to provoke the other side to come to the [negotiating] table." Making this resilient nuclear security complex a reality would require substantial investments in public-private partnerships and improved data infrastructure, particularly to work on classified information at scale. Although the digital twin of the nuclear stockpile would itself be a potential attack surface and would require the greatest possible protection, the case for digital transformation in the nuclear security sector is just as compelling as it is elsewhere.

Opportunity 3: Nuclear Arms Control, Nonproliferation, and Threat Reduction

The United States and the Soviet Union developed a practice of nuclear arms control over decades to promote stability by reducing the incentives for a nuclear first strike or engaging in costly arms racing. That practice involved the creation of what the late Henry Kissinger and Graham Allison refer to as a "conceptual arsenal,[14]" which grew to include a range of practices, including counterforce, secure second strike, and arms control, implemented through negotiated and verified agreements.

Mustafa Suleyman[15], the CEO of Microsoft AI, describes those practices as "the lesson of the Cold War" that "will have to be relearned." No path to technological safety exists without cooperation with adversaries. Although scholars have identified major challenges in applying existing arms control concepts to AI, several interviewees suggested new forms of public-private partnerships to harness emerging technologies for preventing nuclear weapons use and strengthening nuclear arms control, nonproliferation, and threat reduction.

[14] Henry A. Kissinger and Graham Allison, "The Path to AI Arms Control," Foreign Affairs, October 13, 2023 https://www.foreignaffairs.com/united-states/henry-kissinger-path-artificial-intelligence-arms-control.

[15] Mustafa Suleyman with Michael Bhaskar, The Coming Wave: AI, Power, and Our Future (Crown) 2023.

Arms Control

The convergence of emerging technologies can strengthen arms control efforts through enhanced detection, monitoring, and verification technologies (Box 3). As discussed previously, the capability of a wide range of sensors is rapidly advancing, along with the potential to fuse data from large numbers of heterogenous sensors. At the same time, those devices are becoming significantly smaller, lighter, more energy efficient, and more affordable.

Revolutionary changes in sensor and sensor fusion technologies could facilitate new kinds of arms control not previously possible. For example, in their 2023 article "Dual-Use Deception: How Technology Shapes Cooperation in International Relations," Jane Vaynman and Tristan Volpe describe how increased *distinguishability* between civilian and military applications of technology can support successful arms control agreements.[16] An innovation agenda to increase and leverage greater transparency could open new pathways to successful arms control.

Blockchain-enabled tools for enhancing readiness could also provide zero-knowledge proofs to adversaries that the alert status of nuclear forces has not changed. *Zero-knowledge proofs* are cryptographic techniques to prove a fact to another party without revealing details of how that fact is known. Zero-knowledge proofs have long been considered an important objective for verifying the dismantlement of nuclear warheads to prove to another party that an object is or is not a nuclear warhead without revealing classified information. One interviewee emphasized that zero-knowledge proofs could also be applied on a wider scale to verify to adversaries that an entire military base or nuclear arsenal is operating within agreed parameters without divulging more information. Visible Assets Inc., a U.S. company, has a wireless technology called RuBee for "low cost, long-life, low power, wireless, active asset and people tags that work in harsh environments."[17] One interviewee envisioned a future in which RuBee tags could support zero-knowledge proofs, verifying that large numbers of nuclear weapons are in the area where they are declared to be without revealing their specific location.

A congressionally mandated 2023 study by the National Academies of Science, Engineering, and Medicine concluded that the United States "must seek to both modernize MDV [monitoring, detection, and verification] systems and approaches in the near term and revolutionize them in the longer term instead of continuing to make incremental steps forward. Information sharing and fusion of intelligence sources is a key example of an approach with revolutionary potential."[18] The study includes a series of important recommendations representing an innovation agenda for monitoring, detection, and verification technologies.

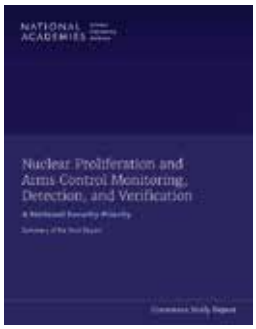

[16] Jane Vaynman and Tristan Volpe, "Dual-Use Deception: How Technology Shapes Cooperation in International Relations," *International Organization* 77, no. 3 (2023): 599–632.

[17] Visible Assets website: https://www.ru-bee.com/VisBlurb/index.html.

[18] National Academies of Sciences, Engineering, and Medicine, *Nuclear Proliferation and Arms Control Monitoring, Detection, and Verification: A National Security Priority: Summary of the Final Report*. (Washington, DC: The National Academies Press., 2023), 1, https://nap.nationalacademies.org/catalog/26558/nuclear-proliferation-and-arms-control-monitoring-detection-and-verification-a-national-security-priority.

BOX 3: **OPEN "SOURCERY": PUBLICLY AVAILABLE INFORMATION AND NUCLEAR SECRETS**

Professor Jeffrey Lewis of the Middlebury Institute for International Studies leads a team pioneering methods of deriving sensitive nuclear security information using commercial technologies, including AI tools and big data. Findings of the team include identifying the location of secret North Korean nuclear weapons facilities.

In 2013, North Korea released a propaganda video that used as a background an image of the interior of a facility for assembly of Transporter-Erector-Launchers (TELs)—specialized vehicles for moving, elevating to a firing position, and launching North Korean nuclear missiles.

Lewis and his team analyzed the video to assess the building, including its dimensions and characteristics, noticing, for example, that the building had unusual window placements along the back, side, and roof. On the basis of that analysis, Lewis's team created three-dimensional (3D) models of the inside and outside of the building using SketchUp, a free, widely available 3D modeling program.

With the models complete, Lewis and the team were able to use open-source information to match their model to locations in the real world. Sources included Google Earth, which has high-resolution satellite and aerial images of areas of North Korea; defector accounts of the locations of North Korean defense industries; and Korean-language social media sites. Once the team narrowed the location to a region used to house defense industries, further satellite images allowed it to locate two plausible buildings that matched the window patterns seen in the video. That effort provided evidence of the location of KN–08 launchers and what appeared to be North Korea's most important facility for the final assembly of TELs.

This example demonstrates how a combination of open-source information (online footage, 3D modeling software, commercially available satellite images, etc.) can give anyone with an Internet connection and keen analytic skills the capability to derive and release the exact location of sensitive facilities that are critical to nuclear weapons supply chains. Moreover, elements of this work that require highly specialized human skills today—including searching images and looking for patterns and deviations—might in the future be automated by computers and LLMs, which has important implications for nuclear security. It demonstrates the potential for open-source research to increase transparency as hiding things becomes increasingly difficult. It shows the increasing ability of civil society to hold governments accountable and to potentially participate in verification of government data and evidence.[19]

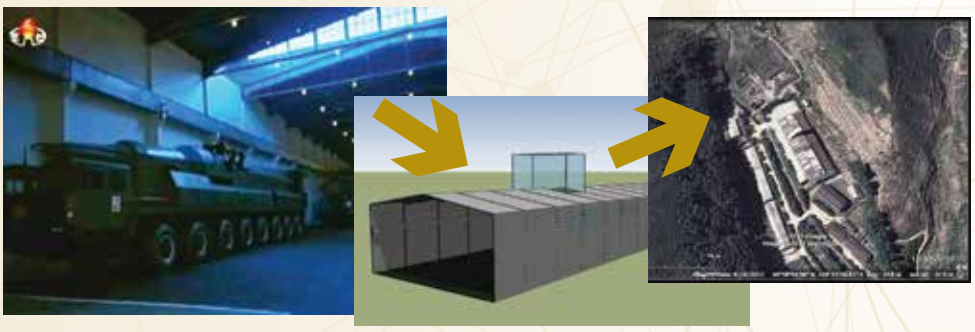

[19] This box was drawn from Jeffrey Lewis, Melissa Hanham, and Amber Lee, "That Ain't My Truck: Where North Korea Assembled Its Chinese Transporter-Erector-Launchers," *38 North*, February 3, 2024, https://www.38north.org/2014/02/jlewis020314/.

Nonproliferation

Nuclear security opportunities may continue to grow as emerging technologies improve transparency and data analysis. Former Department of Homeland Security official Warren Stern observes the emergence of a "revolution in nuclear detection affairs" driven by new kinds of sensors and the use of smartphones interacting with those sensors.[20] The combination of machine learning and open-source information can also contribute to the detection of nuclear proliferation activities (as described in Box 4). One interviewee stressed the benefits of combining AI, ubiquitous sensing, autonomous sensor platforms, and other emerging technologies "using publicly available information with government information and merging space information with trade information—combining data streams that are not robust by themselves but are more robust together."

Threat Reduction

In 1991, the collapse of the Soviet Union left the world's largest nuclear arsenal spread across several newly independent states, many with no experience managing or maintaining a nuclear arsenal. The Nunn-Lugar Cooperative Threat Reduction Program responded by providing support and expertise in crucial tasks, including the deactivation of more than 7,000 nuclear warheads, the destruction of more than 900 ICBMs, and the improvement in securing tons of weapons-usable plutonium and HEU.

The kind of cooperation that was required between Russia, the then-newly independent states, and the Baltic states is difficult to imagine in today's political climate, but the Soviet collapse should chasten us against dismissing the possibility of dramatic political change. If a future political opportunity to resume cooperative threat reduction with a nuclear-armed state appears, a wide range of emerging technological tools not available in the 1990s could contribute to the success of this work.

Much of the site security functionality that required extensive construction in the 1990s can now be achieved with off-the-shelf smart home technologies manufactured by companies such as Roku. For example, modern systems can use acoustic sensors that monitor entire rooms for sounds like breaking glass. Sensors can also notify multiple users across a site in real time via smart devices, rather than a central alarm station that would relay the information by radio, as used in the 1990s.

Handheld light detection and ranging scanning tools, such as those manufactured by the company GeoSLAM, allow real estate agents to quickly create a digital model of a home. This same technology can be used to generate virtual models of nuclear facilities, improving the planning and procurement processes for site security or the movement of weapons and materials. Detailed digital twins of nuclear facilities also can be used to continuously monitor and minimize potential pathways for nuclear-material diversion and to support predictive maintenance of security systems. For example, they can allow inspectors to visualize characteristics such as radiation levels. Such systems also can be used for training purposes and to support verification activities, such as arms-control inspections or International Atomic Energy Agency inspections. AI-driven change detection can add the capacity to predict and monitor metrics of concern for sudden, unexpected, or gradual change that may indicate a situation warranting attention.

[20] Warren Stern, "Revolution in Nuclear Detection Affairs," remarks presented at the American Physical Society and the George Washington University, Washington, DC, November 2-3, 2013 (Brookhaven National Laboratory BNL-103823-2014-CP), https://ia601304.us.archive.org/11/items/RevolutioninNuclearDetectionAffairs/Revolution%20in%20Nuclear%20Detection%20Affairs.pdf.

BOX 4: **NTI-C4ADS STUDY APPLYING MACHINE LEARNING TO NONPROLIFERATION**

For decades, high-risk and illicit trade in nuclear materials, equipment, and technologies has undermined global nuclear nonproliferation efforts. The networks and individuals involved often evade detection by operating within legal systems of trade, finance, transportation, and communication by establishing front companies, forging documents, or laundering money; however, they still leave footprints. Now, an increase in publicly available data provides opportunities to gain visibility into, and expose, such activities. Emerging data science and advanced analytical tools coupled with trade, transport, and other publicly available information (PAI) can and should be used to strengthen global nonproliferation.

From 2019 through 2020, the Nuclear Threat Initiative (NTI) and the Center for Advanced Defense Studies (C4ADS) worked together to demonstrate the power of combining publicly available data with advanced analysis tools to detect high-risk behavior.

By analyzing nuclear trading networks and using machine learning, researchers could identify previously unknown entities of elevated risk within millions of transactions. Automated data preparation using machine learning saved hundreds of analyst hours and identified twice as many potentially high-risk entities as previous manual efforts. In addition, when applied to a baseline study of more than 4 million records, machine-learning techniques helped identify 50 new leads for further review.

Leaders of nonproliferation efforts in governments and multilateral organizations around the world should ensure that the advantages of PAI and modern analytical approaches are applied to monitor and ultimately disrupt illicit nuclear activities. Taken together, these steps will enable the quality, scale, and timeliness needed for regional and, perhaps in the future, global monitoring.

Keeping a "human in the loop" is just the first necessary step in adjusting U.S. nuclear deterrence policy and posture to the astonishing pace of AI and other emerging technologies.

Conclusions, Recommendations, and Next Steps

This report was undertaken to identify ways that AI and other technologies emerging from the commercial sector could present risks and opportunities for nuclear security that are not adequately addressed by keeping a "human in the loop" of nuclear command and control. In the course of 32 interviews with leading experts, the study team identified three significant categories of risk:

- Novel and newly expanded risks to nuclear forces
- New pathways to nuclear escalation
- Increased risk of nuclear proliferation and terrorism

And the study team identified three opportunities:

- Improved warning confidence
- Deterrence resilience
- New pathways to cooperative security in areas including nonproliferation, arms control, threat reduction, and nuclear disarmament

Nuclear modernization focused on the replacement of Cold War capabilities may not adequately respond to the emerging technological environment in which extremely rapid commercial innovation at scale is interwoven with questions of nuclear security. Neither the private sector nor governments are working toward a comprehensive understanding of the scale and effects of this technological revolution in nuclear security. Relatively little attention has been paid to the possibility that the convergence of emerging technologies may pose new risks to nuclear security by challenging the systems and approaches that nuclear-armed states use to prevent nuclear weapons use and to "control" nuclear weapons. Ultimately, the success of governments in mitigating the nuclear security risks and reaping the nuclear security benefits of this technological revolution will depend on leadership that can move beyond traditional approaches.

The United States is on course to spend $1.2–$1.7 trillion reflexively to replace nuclear weapons and delivery platforms designed for the 20th century without carefully considering what nuclear war prevention requires now. The United States should change course so that it develops an updated strategy to prevent nuclear weapons use and to reduce the dangers of nuclear proliferation and nuclear terrorism by leveraging emerging technologies in maximally advantageous ways to build a new nuclear security order.

Nuclear security will be a moving target far into the future, with ever-changing tools, strategies, and technological requirements. Humanity needs a global innovation ecosystem to meet this continuous challenge now and into the future, and creating a national innovation ecosystem for preventing nuclear war in the United States would be an important first step. The authors of this report offer the following recommendations toward this objective:

Recommendation 1: Establish a National Security Commission on Nuclear Security Innovation

The National Security Commission on Artificial Intelligence (NSCAI), chaired by former Google CEO Eric Schmidt and former Deputy Secretary of Defense Bob Work, made a crucial contribution to understanding how AI will affect nuclear security. NSCAI's March 2021 final report specifically recommended that the United States "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." As a result, the 2022 Nuclear Posture Review contained the following statement: "In all cases, the United States will maintain a human 'in the loop' for all actions critical to informing and executing decisions by the president to initiate and terminate nuclear weapon employment"; this position was a sharp improvement over the 2018 Nuclear Posture Review, which did not mention artificial intelligence at all. But keeping a "human in the loop" is just the first necessary step in adjusting U.S. nuclear deterrence policy and posture to the astonishing pace of AI and other emerging technologies. A continuous approach is needed (a) to identify and manage the spectrum of risks and seize the most promising opportunities that emerging technologies pose for nuclear security and (b) to foster the adoption of stabilizing approaches to emerging technology in nuclear security internationally. In the National Defense Authorization Act for fiscal year 2022, the U.S. Congress required the secretary of defense to conduct a "nuclear fail-safe review" assessing the safety, security, and reliability of nuclear weapons and related systems. This classified technical study has been completed and should be an input to further work to understand how the convergence of emerging technologies affects U.S. nuclear strategy, policy, acquisitions, force structure, and posture. We recommend that the U.S. Congress should establish a commission on nuclear security in the age of AI modeled on the National Security Commission on AI.

- The commission should be charged with producing a comprehensive report with recommendations for nuclear security technology development, diplomatic engagement, and public-private partnerships.
- The commission should specifically address the human layer and network vulnerabilities of nuclear security systems.
- The commission should comprise leaders in nuclear security and various branches of emerging technology that could accelerate appropriate responses to nuclear security dangers. Congress should provide appropriate resources and staff to this commission.

Recommendation 2: Create Incentives for Transformational Engagement with the Private Sector in Nuclear Security Innovation

The U.S. government should create structures and incentives for deep public-private partnerships for nuclear security, such as a nuclear security innovation unit to accelerate adoption of leading commercial technology in the U.S. nuclear security complex using the model of the Defense Innovation Unit; a nuclear security accelerator to quickly deliver nuclear security capabilities based on AI and emerging technologies into the nuclear security mission space modeled on the Space Development Agency and its motto "Semper Citius"—"Always Faster"; and a "NucWERX" based on AFWERX and SpaceWERX to accelerate agile and affordable capability transitions by teaming leaders in innovative technology with the U.S. national laboratories to forge an innovation ecosystem that delivers disruptive nuclear security capabilities.

Congress should also look beyond those models to types of opportunities that leverage the unique innovation capabilities of the U.S. national laboratory system through place-based innovation, multi-laboratory collaboration, and the establishment of new national laboratories as appropriate.

Next Steps

NTI can build on this work by convening an expert study group on emerging technologies and nuclear security to test and address the risks and opportunities identified in this report over the next year. Concluding the work of that group, NTI can engage governments to design and develop a global innovation ecosystem for nuclear war prevention.

Appendix: Interviewees

James Acton, PhD
Co-director, Nuclear Policy Program
Carnegie Endowment for International Peace

Yasir Atalan, PhD
Data Fellow, Futures Lab
Center for Strategic and International Studies

Kim Budil, PhD
Director
Lawrence Livermore National Laboratory

Lyndon Burford, PhD
Policy Fellow and Project Manager, Nuclear Transparency Index
BASIC

Lord Browne of Ladyton
Vice Chair
Nuclear Threat Initiative

Cory Doctorow
Author and Journalist

Thomas Fanning
Former Executive Chairman
Southern Company

Peter Fisher, PhD
Thomas A. Frank Professor of Physics (1977)
Massachusetts Institute of Technology

James Gosler
Senior Fellow, Applied Physics Laboratory
Johns Hopkins University

Nola Haynes, PhD
Adjunct Professor
Georgetown University

Martin Hellman, PhD
Professor Emeritus of Electrical Engineering
Stanford University

Jill Hruby
Former Under Secretary of Energy for Nuclear Security, National Nuclear Security Administration
Department of Energy

Shirley Ann Jackson, PhD
President Emerita
Rensselaer Polytechnic Institute

Scott Kemp, PhD
Director, MIT Laboratory for Nuclear Security and Policy
Massachusetts Institute of Technology

Dimitri Kusnezov, PhD
Former Under Secretary for Science and Technology
Department of Homeland Security

LTC Lincoln Leibner (ret.)
Director, Special Projects
U.S. Fleet Cyber Command/U.S. Tenth Fleet

Jeffrey Lewis, PhD
Distinguished Scholar of Global Security
Middlebury College

Patricia Lewis, PhD
International Security Expert

Herbert Lin, ScD
Senior Research Scholar for Cyber Policy and Security, Center for International Security and Cooperation
Stanford University

Ernest Moniz, PhD
Co-Chair and Chief Executive Officer
Nuclear Threat Initiative

Adm. Mike Mullen (ret.)
17th Chair, Joint Chiefs of Staff
United States Navy

Chantell Murphy, PhD
Program Manager, Nonproliferation and Arms Control Research and Development Group
Y–12 National Security Complex

Lindsay Rand, PhD
Research Scholar
U.C. Berkeley Risk and Security Lab

Ray Rothrock
Venture Investor
FiftySix Investments

Alice Santini
AI-Nuclear Policy Advisor
Institute for Security and Technology

Nathalie Tocci, PhD
Director
Istituto Affari Internazionali

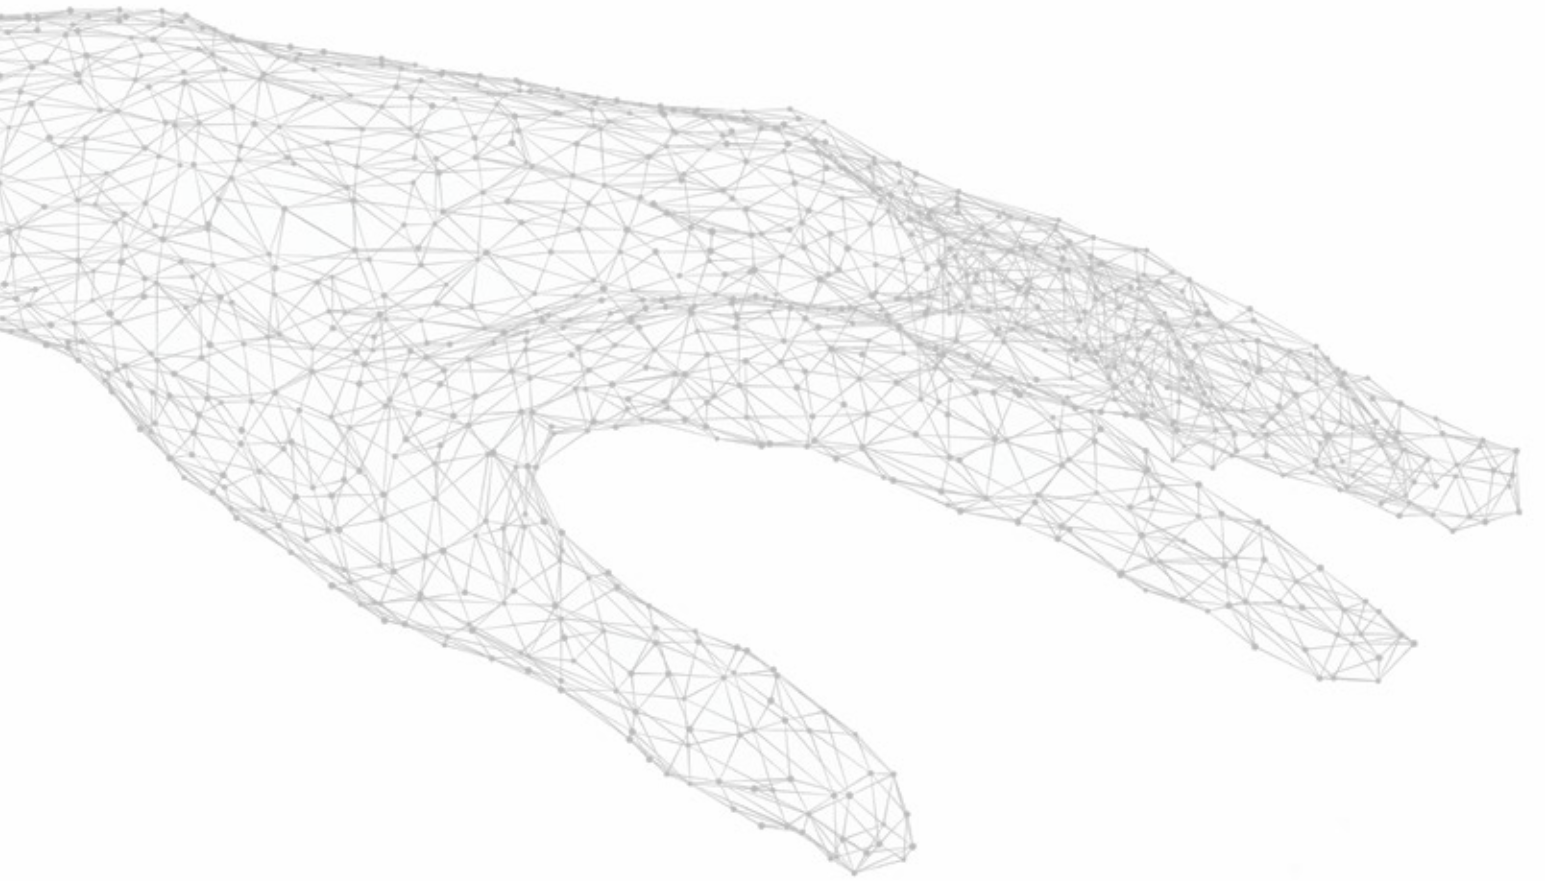

Jane Vaynman, PhD
Assistant Professor, School of Advanced International Studies
Johns Hopkins University

Cindy Vestergaard, PhD
Project Lead, Converging Technologies and Global Security Program
Stimson Center

Nina Wagner
Principal Director, Nuclear & Countering WMD Policy
Office of the Secretary of Defense

Adm. James "Sandy" Winnefeld (ret.)
9th Vice Chair of the Joint Chiefs of Staff
United States Navy

Tong Zhao, Ph.D.
Senior Fellow, Carnegie China and Nuclear Policy Program
Carnegie Endowment for International Peace

Maria Zuber, PhD
E.A. Griswold Professor of Geophysics and Presidential Advisor for Science and Technology Policy
Massachusetts Institute for Technology

Glossary

ARTIFICIAL INTELLIGENCE (AI)
Software that enables computers to make decisions historically reserved for humans

AGENTIC ARTIFICIAL INTELLIGENCE
A class of artificial intelligence designed to act with autonomy, able to conduct strategic planning and dynamic problem solving without direct human intervention

ATTACK SURFACE
A group of paths, methods, or scenarios that can be used to enter data to, extract data from, or control a device or software in an environment

BEHAVIORAL NUDGE
Gentle intervention designed to influence people's decisions and steer individuals toward specific outcomes without limiting their freedom of choice

DATA LAKE
A centralized repository that stores large volumes of data in its original form

DEEPFAKE
An image or recording that has been convincingly altered to misrepresent someone as doing or saying something

DIGITAL EXHAUST
An invisible trail of data left behind by a person's interactions with technological services

DIGITAL TWIN
Virtual model of an intended or real-world object for purposes such as simulation, monitoring, or maintenance

DISTRIBUTED LEDGER TECHNOLOGY (DLT)
A system in which data are stored and synchronized over several geographical locations, as opposed to a central database, and that does not require a central administrator

FULL-SPECTRUM CYBER
The combined arms employment of joint military capabilities not only to exploit but also create advantages in the cyber domain

HALLUCINATION
A response made by an AI that contains false or misleading information presented as fact

HUMAN LAYER
The component of a system comprising humans as opposed to the hardware or software

INTERNET OF THINGS
A network of physical devices, vehicles, appliances, and other objects embedded with sensors, software, and network connectivity, enabling them to collect and share data

LARGE LANGUAGE MODEL (LLM)
A type of machine learning model tasked with natural language processing, such as language generation

MACHINE LEARNING
A form of AI that uses statistical algorithms that learn from data and performs tasks without explicit instructions

NONPROLIFERATION
The practice of preventing the spread of nuclear weapons, fissionable materials, and weapons-applicable nuclear technology

NUCLEAR ARSENAL/ENTERPRISE/STOCKPILE/COMPLEX
The body of infrastructure and expertise dedicated to the development, testing, and maintenance of nuclear weapons

NUCLEAR DETERRENCE
The threat or implied threat of nuclear use against an adversary to shape that adversary's behavior

SUPPLY CHAIN
A logistics system that converts raw materials into completed products and distributes them to end users

SYNTHETIC DATA
Data that is artificially generated rather than created by real-world events, typically used for training AI models

About the Authors

Douglas Shaw, PhD

Senior Advisor, FutureSafe: AI and Emerging Technology

Before joining NTI, Douglas Shaw served in a variety of senior leadership roles in higher education, including as senior associate provost at George Washington University, associate dean of the Elliott School of International Affairs, and director of policy planning at Georgetown University. During the Clinton administration, he served in the U.S. Department of Energy, working to secure fissile materials in Ukraine, and at the U.S. Arms Control and Disarmament Agency, working to indefinitely extend the Treaty on the Non-Proliferation of Nuclear Weapons.

Shaw holds BSFS, MA, and PhD degrees from Georgetown University in international relations and security studies. He taught international security studies and nuclear weapons policy at George Washington University and Georgetown University for two decades.

Isabelle Williams

Senior Director, Global Nuclear Policy Program

Isabelle Williams leads and supports NTI projects related to building and sustaining a world without nuclear weapons, reassessing nuclear deterrence, and promoting the understanding of cascading nuclear effects. Williams also supports NTI-wide implementation of institutional objectives and strategic priorities and serves as NTI liaison for Horizon 2045 (a collaborative effort to address the interconnectedness of global threats, including nuclear weapons).

Williams holds a BA and an MA in international studies from the University of Leeds, United Kingdom. She previously held positions with the Partnership for Global Security, the Chemical and Biological Arms Control Institute, and the International Institute for Strategic Studies.

Patricia Jaworek

Director, Global Nuclear Policy Program

Patricia Jaworek supports NTI's efforts to reduce global nuclear risks, focusing on U.S.-Russia arms control, Euro-Atlantic security, and the nuclear Non-Proliferation Treaty. She leads NTI's activities to advance understanding of the global and long-term effects of nuclear weapons use and their implications for nuclear policy. Jaworek previously worked as a research assistant at the NATO Parliamentary Assembly in Brussels, where she focused on NATO deterrence and defense policy.

Jaworek holds a joint master's degree in transatlantic affairs from the Fletcher School of Law and Diplomacy at Tufts University and the College of Europe and a law degree from the University of Hamburg, with a specialization in European and public international law.

Kevin Park

Former Program Associate, FutureSafe: AI and Emerging Technology

Kevin Park supported the FutureSafe program's work to identify the nuclear security policy implications of disruptive commercial technologies. Previously, he served as an infantryman in the Republic of Korea Marine Corps.

Park is a recent graduate of George Mason University's Master of International Security program. Park studied and researched numerous national security topics, including nuclear arms control, maritime security, and the impact of emergent technologies on international security.

Pravin Rajan

Pravin Rajan is a quantitative trader. He was previously the founder of FreemarketsAI, a venture-backed startup that applied machine learning to markets. Before that, he served as a Marine in Afghanistan and at the Pentagon, where he conducted an analysis on American security interests over the next 30 years.

Rajan has a BS in Science, Technology, and International Affairs from Georgetown's School of Foreign Service and an MPhil in Social Anthropology from Oxford University, where he was a Rhodes Scholar.

1776 EYE STREET, NW, SUITE 1000
WASHINGTON, DC 20006
(202) 296–4810
WWW.NTI.ORG


NTI.ORG


NTI_WMD


NTI_WMD


NUCLEAR THREAT INITIATIVE