




AI Is Already at War

How Artificial Intelligence Will Transform the Military

By **Michèle A. Flournoy** November/December 2023

Published on October 24, 2023



Eduardo Morciano



34:47
Play

In 2002, a special operations team practiced raiding a safehouse. The team silently approached a two-story building, built for military training, where a fictitious terrorist leader was hiding. One soldier crept up to an open window and tossed in a small drone piloted by artificial intelligence. The AI drone began flying autonomously through the

building, room by room, beaming footage from its camera directly to the commander's handheld tablet outside. In just a few minutes, the team had full situational awareness of the interior of the building. It knew which rooms were empty, which were occupied by sleeping family members, and where the primary target was. The team entered the building knowing exactly where to go, reducing the risk for each member. The drill was a success: had it been real, the team would have killed the terrorist leader.

The AI-piloted quadcopter, designed by Shield AI (where I was an adviser), has since been used in real-world operations. It is just one of the many ways that AI is beginning to reshape U.S. national security. The U.S. military is using AI to optimize everything from equipment maintenance to budgetary decisions. Intelligence analysts are relying on AI to quickly scan mountains of information to identify relevant patterns that enable them to make better judgments and to make them faster. In the future, Americans can expect AI to change how the United States and its adversaries fight on the battlefield, as well. In short, AI has sparked a security revolution—one that is just starting to unfold.

As AI has burst into the public consciousness, some researchers, worried about AI's dangers, have called for a pause on development. But stopping American AI progress is impossible: the mathematical foundations of AI are ubiquitous, the human skills to create AI models have widely proliferated, and the drivers of AI research and development—both human creativity and commercial gain—are very powerful. Trying to stop progress would also be a mistake. China is working hard to surpass the United States in AI, particularly when it comes to military applications. If it succeeds, Beijing would then possess a much more powerful military, one potentially able to increase the tempo and effect of its operations beyond what the United States can match. China's ability to use cyber and electronic warfare against U.S. networks and critical infrastructure would also be dangerously enhanced. Put simply, the Pentagon needs to accelerate—not slow—its adoption of responsible AI. If it doesn't, Washington could lose the military superiority that underwrites the interests of the United States, the security of its allies and partners, and the rules-based international order.

Stay informed.

In-depth analysis delivered weekly.

[Sign Up](#)

Acceleration, however, is easier said than done. The United States may lead the world when it comes to artificial intelligence research and development, but the U.S. government still struggles to adopt innovative technologies such as AI with speed and at scale. It does not employ enough professionals with the technical expertise needed to test, evaluate, procure, and manage AI products. It is still building the data and computer infrastructure necessary to support large AI models. It lacks the flexible funding required to quickly take the most promising AI prototypes and scale them across agencies. And it has yet to build up the testing and evaluation processes and platforms needed to ensure that any AI integrated into military systems is safe, secure, and trusted. When AI plays a role in the use of force, the bar for safety and reliability must remain very high.

Politicians and defense officials are aware of these issues. Congressional leaders are paying close attention to AI, and they are discussing how they can regulate the industry and yet keep it globally competitive. The Office of the Secretary of Defense has issued a policy framework for AI to expedite its responsible and safe adoption by the Defense Department. The essential effort to simultaneously foster AI and put guardrails around its use—aims that are seemingly in tension—is underway.

But Congress has yet to act, and the implementation of the Pentagon's AI framework is still very much a work in progress. Although the creation of a Chief Digital and Artificial Intelligence Office at the Defense Department was an important milestone, Congress has yet to provide this office with the resources it needs to drive responsible AI adoption across the defense establishment. To ensure that AI defense applications are both safe and successful, the Pentagon will need to further bolster AI guardrails, add new technical staff, and develop new ways of testing and procuring AI. Time is of the essence, and the stakes are too high for the United States to fall behind.

HERE AND NOW

Even as policies and regulations are still being written, AI is already transforming U.S. security. The U.S. Air Force, for example, is beginning to use AI to help it allocate resources and to predict how a single decision can reshape its program and budget. If air force leaders, for example, add another squadron of F-35s, their AI-enabled resource allocation platform can immediately highlight not only the direct costs of the decision but also its effects on personnel, bases, aircraft availability, and other important domains.

Similarly, the military is beginning to use AI models in the maintenance of complex weapons systems, from ships to fighter jets. AI programs can now collect data from a platform's sensors and predict when and what kind of maintenance will maximize its readiness and longevity while minimizing costs.

These maintenance insights are tremendously helpful, and they are just the beginning of what predictive AI can do. The U.S. intelligence community and several U.S. combatant commands—the joint military commands with operational responsibility for a particular region or function—are using AI to sift through reams of classified and unclassified data to identify patterns of behavior and forecast future international events. In the intelligence community, AI helped analysts predict Russia's invasion of Ukraine months in advance, enabling the United States to warn the world and deny Russian President Vladimir Putin the element of surprise. At U.S. Strategic Command, AI developed by Rhombus Power (where I am an adviser) is being used to help warn officials about the movement of nuclear-armed missiles that often evaded detection in the past.

Predictive AI could also give Washington a better understanding of what its potential adversaries might be thinking, especially leaders in Beijing. Unlike during the height of the Cold War, when there were legions of experts on Soviet decision-making, the United States is still figuring out how China's leadership translates policy into specific actions. The intelligence community could, for instance, develop a large language model that would ingest all available writings and speeches by Chinese leaders, as well as U.S. intelligence reports about these figures, and then emulate how Chinese President Xi Jinping might

decide to execute stated policy. Analysts could ask the model specific questions—“Under what circumstances would President Xi be willing to use force against Taiwan?”—and anticipate potential responses based on a wealth of data from more sources than any human being could ever quickly synthesize. They could even ask the model to map out how a crisis might unfold and how different decisions would shape the outcome. The resulting insights could be useful in informing analysts and policymakers, provided the training sets were transparent (meaning they cite the sources of data underlying key judgments) and trusted (not prone to “hallucinations”—inexplicable inferences made by AI).

AI has sparked a security revolution that is just starting to unfold.

Intelligence officers are already using AI daily to sift through thousands of pictures and videos. In the past, analysts had to watch thousands of hours of full-motion video to find and tag objects of interest, whether a concentration of tanks or dispersed mobile missiles. But with AI, developers can train a model to examine all this material and identify only the objects the analyst is looking for—usually in a matter of seconds or minutes. The analyst can also set the AI model to send an alert whenever a new object of interest is found in a given geographic area. These “computer vision” tools enable analysts to spend more time doing what only humans can do: applying their expertise and judgment to assess the meaning and implications of what AI discovers. As these models become more accurate and trusted, they have the potential to help U.S. commanders on the ground make critical operational decisions much faster than an adversary can respond, giving U.S. forces a tremendous—perhaps even decisive—advantage.

AI could support military operations in other ways, as well. For instance, if an adversary were to jam or attack U.S. command, control, and communications networks, AI could enable a smart switching and routing agent that would redirect the flow of information between sensors, decision-makers, and shooters to make sure they stay connected and can maintain situational awareness. Having these capabilities will be critical to ensuring that Washington and its allies can make better decisions faster than their adversaries, even in the thick of combat.

AI could further help U.S. and allied forces by amplifying the work of individual service members in the field. Some AI applications currently in development allow a single human operator to control multiple unmanned systems, such as a swarm of drones in the air, on the water, or undersea. For example, a fighter pilot could use a swarm of flying drones to confuse or overwhelm an adversary's radar and air defense system. A submarine commander could use undersea unmanned vehicles to conduct reconnaissance in a heavily defended area or to hunt for undersea mines that threaten U.S. and allied ships. The Pentagon recently announced its Replicator drone program, which promises to field thousands of small, smart, low-cost, expendable, autonomous systems within the next two years.

In a conflict with China over Taiwan, this human-machine teaming could be critical. If Beijing decides to use force to claim the island, China will have the advantage of fighting in its own backyard, allowing it to mass forces more easily. The United States, meanwhile, will be sending its units long distances and in far fewer numbers. If the U.S. military can augment its manned platforms such as fighters, bombers, ships, and submarines with large numbers of relatively cheap unmanned systems, it could compensate somewhat for this comparative disadvantage and greatly complicate the Chinese military's operations.

PLAY IT RIGHT

Beijing, of course, has no intention of ceding technological dominance to Washington. It is working hard to develop its own advanced AI military applications. China is investing heavily in many of the same AI use cases as the United States—such as surveillance, target identification, and drone swarms. The difference is that it may not be bound by the same ethical constraints as the United States and its allies, particularly when it comes to using fully autonomous weapons systems.

In the race for technological supremacy, China has some obvious advantages. Unlike Washington, Beijing can dictate its country's economic priorities and allocate whatever resources it deems necessary to meet AI targets. China's national security policy encourages Chinese hackers, officials, and employees to steal Western intellectual property, and Beijing

is unabashed in trying to recruit leading Western technologists to work with Chinese institutions. Because China has a policy of “civil-military fusion,” which eliminates barriers between its civilian and military sectors, the People’s Liberation Army can draw on the work of Chinese experts and companies whenever it likes. And by 2025, China will churn out nearly twice as many Ph.D. candidates in science, technology, engineering, and mathematics as the United States does, flooding China’s economy with talented computer scientists in particular.

But the United States has its own unique strengths. The country’s market-based economy and more open political system give developers room to be creative. It has unrivaled innovation ecosystems in Silicon Valley, the Austin metropolitan area, the Massachusetts Route 128 corridor, and elsewhere. The United States also has a vibrant venture capital and private equity ecosystem that draws incomparable domestic and international investment. It is home to many of the world’s leading universities, allowing it to attract and retain some of the world’s best tech talent. Indeed, half the startups in Silicon Valley have at least one founder who is an immigrant. Even among those who lament China’s rapid AI progress, few, if any, would trade the United States’ hand for China’s. But almost all of them would agree the United States needs to play its hand better to win.

To do so, the Defense Department and the intelligence community will have to invest more in accelerating AI adoption. They can start by building common digital infrastructure systems that share the same standards to ensure interoperability. The infrastructure would include cloud-based technologies and services; common data standards; validated data sets; shared access to secure software stacks; sophisticated tools for the testing, evaluation, and validation of AI models; and secure application programming interfaces that control who gets access to what information at various levels of classification. The goal would be to give developers the data, algorithms, tools, and compute power—or high-speed computing power—they need to create, test, validate, and use new AI tools.



A U.S. soldier wearing a data display system, Camp Pendleton, California, March 2018

Rhita Daniel / U.S. Marine Corps / Reuters

Those tools will only be as good as the people who operate them, of course, and right now, the Defense Department does not have a digitally adept workforce. Few people on staff understand enough about AI to properly govern its use, to test and evaluate AI tools to ensure they meet the Pentagon’s “responsible AI” standards, or to assess which AI models best meet the needs of the military or the Defense Department—one of the world’s largest enterprises.

To attract more AI talent and to make better use of the tech workforce it already has, the Defense Department will need to improve how it recruits and manages digitally skilled employees. The Pentagon can start by following the advice of the National Security Commission on AI and establish a digital corps (modeled on the Army Medical Corps) to organize, train, and equip technologists. In addition, all the existing military service academies should start teaching the basics of AI, and the Pentagon should also establish a U.S. digital service academy that would educate and train aspiring civilian technologists,

offering them a free college education in exchange for a commitment to serve in government for at least five years after graduating. Finally, the Defense Department should create a digital reserve corps in which tech workers from across the United States could volunteer, part time, to serve their country.

The Pentagon, however, will never be able to attract as many AI experts as the private sector. The defense establishment must therefore improve how it leverages outside talent. For starters, the Defense Department should deepen its conversations with technology companies and the computer science departments of leading universities. It should also reduce some of the outdated barriers to tech firms doing business with the government. To do so, defense officials must rethink how they buy software-based products and services, including AI. Instead of taking years to develop a fixed set of highly specific requirements—as the department does when procuring military hardware—it should quickly identify the specific problems it is trying to solve and the common standards that any proposed solutions must meet and then allow companies to offer solutions in a competitive bidding process. It should also make sure that the people who will actually use the specific AI tools are able to provide feedback as models are being developed and tested.

In fact, the Pentagon should create a dedicated career path for acquisition professionals who want to specialize in AI and other commercially driven technologies. Most of the Defense Department's current acquisition corps have been trained to buy complex weapons systems, such as submarines, missiles, and jets, which requires paying meticulous attention to whether contractors meet rigid specifications, cost requirements, and scheduled milestones. As a result, most of these professionals are (understandably) highly risk averse—they are neither trained nor incentivized to buy rapidly developing commercial technologies or to disrupt an existing multiyear acquisition program to integrate a more effective new technology. The Pentagon should therefore create a new cohort of acquisition experts who are specifically trained to buy these kinds of systems. This cadre should be considered the Green Berets of the acquisition force, and its members should be rewarded and promoted based on their ability to quickly deliver and scale needed commercial technologies, such as AI.

Beijing has no intention of ceding technological dominance to Washington.

Although internal reforms will help the Pentagon accelerate progress, defense officials will also need sustained congressional support to keep pace with their Chinese counterparts. To that end, Congress should give the Defense Department more flexible funding that allows it to optimally manage AI programs. Most of the Pentagon's appropriations are fixed: when Congress funds a program, the department cannot simply redirect the money to something else. But AI is evolving so fast, and in so many different directions, that defense officials need more reprogramming authorities and more flexible funding so they can quickly move money out of underperforming projects and reinvest it in more promising ones, giving Congress appropriate notice. This approach is critical to enabling the Pentagon to adopt AI with more agility and speed.

Congress should simultaneously provide the Chief Digital and Artificial Intelligence Office with bridge funding to help promising AI pilot projects cross the so-called valley of death—the difficult period between when a project demonstrates success and when the department is ready to make it a full-scale program of record. The U.S. military simply cannot afford to delay the adoption of a critical AI tool that emerges in 2023 until the 2025 budget or later.

The United States will also need to continue attracting the best tech talent in the world, including by reforming elements of the U.S. immigration system. Science and technology students and workers may want to come to and stay in the United States, but byzantine immigration rules make it impossible for many of them to do so. Educational visas, for instance, do not let foreign students stay in the United States for more than three years after graduation. The resulting dynamic is perverse: U.S. institutions train many of the world's best tech experts, only to send them away. Many of them are Chinese and return to China.

In addition, congressionally imposed caps on H-1B visas—the visa the United States most commonly offers to skilled workers—mean that the country can bring in only a small percentage of people who apply. For example, from the 758,994 eligible electronic registrations received during the 2023 H-1B lottery, only 110,791 people were selected (or less than 15 percent). In short, the United States is keeping out much-needed foreign talent

that would willingly and meaningfully contribute to the country's ability to compete in AI and other critical technologies.

HIGH RISK, HIGH REWARD

AI is indispensable to the United States' future security. But it also poses major risks. AI is already accelerating the spread of disinformation online and facilitating inadvertent discrimination in hiring. Computer scientists have argued that it could enable automated cyberattacks at "machine speeds," as well. Chemists have shown that AI can synthesize chemical weapons, and biologists have expressed concern that it could be used to design new pathogens or bioweapons. The risks are severe enough that even AI industry leaders have expressed alarm. In May 2023, the heads of almost every major U.S. AI lab signed a letter warning that their inventions could pose an existential threat to humanity.

Indeed, national security is the realm of human activity where the risks of AI are most profound. AI models could, for example, misidentify people or objects as targets, resulting in unintended death and destruction during conflict. Black box AI models—ones whose reasoning cannot be adequately understood or explained—might lead military planners to make hazardous decisions. This risk would be most acute if AI developed for one situation were applied to another without enough testing and oversight. What might be perfectly rational and responsible in one situation might be irrational and dangerous in another.

The risks do not stem just from poorly designed or carelessly used systems. The United States could be fastidious in developing and implementing AI, only for its adversaries to find ways to corrupt U.S. data, prompting systems to go haywire. For example, if an adversary were able to spoof an AI-enabled computer vision tool into targeting a civilian vehicle instead of a military one, it could cause the United States to inadvertently harm civilians in a conflict zone, undermining U.S. credibility and moral authority. An adversary could also corrupt data in ways that would degrade the performance of an AI-enabled weapon system or that could cause it to shut down.

The Pentagon is aware of these risks, and in February 2020, it issued a set of ethical

principles governing how AI should be used. One principle called on the department's personnel to exercise judgment and care in developing, deploying, and using AI capabilities. Another said the Defense Department will try to "minimize unintended bias in AI capabilities." A third called for ensuring that all AI is made and used in ways that can be understood and explained—with data and methodologies that are transparent and auditable. And defense leaders have directed their employees to make sure that AI systems are rigorously tested for their safety, security, and effectiveness; that AI systems are assigned to clearly defined uses; and that AI systems can be disengaged or deactivated if they exhibit unintended behavior.



A drone flying above Ein Shemer, Israel, July 2023

Amir Cohen / Reuters

For autonomous and semiautonomous weapons, the Defense Department has issued even more specific guidance. Pentagon leaders have directed commanders and operators to use careful judgment over AI-enabled weapons, including by ensuring that these weapons are used in ways that are consistent with the parameters of the model's training and with the

rules of engagement for the operation in which the AI is being deployed. The Defense Department's rules also stipulate that commanders use AI in accordance with the laws of war. For example, any AI-enabled weapon must be discriminate, able to distinguish between combatants and noncombatants on the battlefield, and able to avoid deliberately targeting the latter. The Pentagon has also forbidden the use of AI in its nuclear command-and-control systems, and it has urged other nuclear powers to do the same.

Among the U.S. defense community's leadership, these "responsible AI" rules have achieved great consensus. But putting them into practice is no small challenge—especially given the size of the United States' defense apparatus. The Pentagon has started the process by creating a high-level governance body, beginning to establish data and digital infrastructure to support a variety of AI applications; building out the testing, evaluation, and validation capabilities needed to ensure compliance with the Defense Department's AI principles; and increasing AI awareness across the department. This implementation process is still in its infancy. But the policy framework provides a sound basis on which to build.

Still, the Pentagon would be wise to further strengthen these guidelines. For example, defense officials should require that all AI vendors give the Defense Department full transparency into the origins of data they use in their training sets. In addition, the department should make sure that the behavior of any AI model it adopts is explainable (fully understood by its users and developers), without stifling innovation. It can do so by strengthening how it tests, evaluates, and verifies systems. The department should also scale and broaden the work done by the Defense Advanced Research Projects Agency—one of the entities responsible for developing emerging technologies for the military—on making sure that AI tools are explainable and responsible by design. The department's ethical principles, in other words, should be treated as required traits that shape how defense AI models are designed from the start.

U.S. institutions train many of the world's best tech experts, only to send them away.

But the U.S. defense community will not be able to speed AI adoption unless the public believes it will use AI in ways that are effective, responsible, ethical, and lawful. Otherwise, the first time an AI application leads to a very bad decision or serious unintended consequences on the battlefield, warfighters are unlikely to trust it, and policymakers and lawmakers are likely to suspend or prohibit its use. The Defense Department must therefore increase its investment in the research and development of AI safety and security. It must be transparent about what it will and will not use AI to do. And the Pentagon should consider making its vendors put guardrails on how they develop AI. If a company wants to provide AI to the military, for example, the Defense Department could require it to meet rigorous data protection and cybersecurity standards. By doing so, the Pentagon could help make AI safer, not just for the armed forces, but for everyone.

The United States, of course, cannot singlehandedly make sure that AI is developed and used responsibly. Other countries—including competitors—will also have to adopt policy guardrails and norms. The world took a valuable first step when, in November 2021, 193 countries approved a global agreement on the ethics of artificial intelligence—the world's first. It includes the principle that countries must guarantee human oversight of and agency over all AI.

Although this agreement is an important foundation, the United States should seek out venues to discuss AI with its potential adversaries, especially China, just as it found ways to talk about nuclear weapons and other forms of arms control with the Soviet Union during the Cold War. To succeed, Washington will also have to work closely with its allies and partners to make sure they are all on the same page. Countries that agree on a set of AI norms should be willing to threaten violators with severe costs, including multilateral economic sanctions, expulsion from international forums, and legal action to hold perpetrators responsible for damage. Actors that violate AI rules, for instance, could be indicted in a U.S. federal court, as five Chinese hackers were in 2014 for launching cyberattacks on U.S. companies. States that violate these rules could face potential retaliation for any harm done—including, in extreme cases, military action.

THE NEED FOR RESPONSIBLE SPEED

In the world of microelectronics, experts often talk about Moore's law: the principle that the number of transistors on chips doubles every two years, resulting in exponentially more capable devices. The law helps explain the rapid rise of so many technological innovations, including smartphones and search engines.

Within national security, AI progress has created another kind of Moore's law. Whichever military first masters organizing, incorporating, and institutionalizing the use of data and AI into its operations in the coming years will reap exponential advances, giving it remarkable advantages over its foes. The first adopter of AI at scale is likely to have a faster decision cycle and better information on which to base decisions. Its networks are likely to be more resilient when under attack, preserving its ability to maintain situational awareness, defend its forces, engage targets effectively, and protect the integrity of its command, control, and communications. It will also be able to control swarms of unmanned systems in the air, on the water, and under the sea to confuse and overwhelm an adversary. The United States cannot afford to fall behind.

But the national security apparatus cannot afford to be reckless, either. Without proper safeguards, AI models could cause all kinds of unintended harm. Rogue systems could even kill U.S. troops or unarmed civilians in or near areas of combat. The United States therefore finds itself facing a conundrum. The stakes of slowing AI down are unacceptably high, but so are the stakes of racing ahead without needed precautions.

U.S. policymakers appear to understand this paradox. Congressional leaders know that if they were to regulate AI with too heavy a hand, they could prompt the best AI innovators to leave the United States to work where there are fewer restrictions, and the United States would then fall behind its competitors. But both Democratic and Republican policymakers also know that some regulation and oversight is essential to ensuring that AI adoption is safe and responsible. The House of Representatives and the Senate are holding sessions to educate their members and scheduling hearings to get advice from experts. These efforts to build bipartisan consensus before legislating should be applauded.