ISSUE 102

STRATEGIKA

Conflicts Of The Past As Lessons For The Present

NOVEMBER 2025

THE APPLICATION OF ARTIFICIAL INTELLIGENCE TO MODERN WEAPONRY

IN THIS ISSUE

BRADLEY L. BOYD • ARIEL CONN • LENA TRABUCCO

CONTENTS

November 2025 • Issue 102

ABOUT THE POSTERS IN THIS ISSUE

Documenting the wartime viewpoints and diverse political sentiments of the twentieth century, the Hoover Institution Library & Archives Poster Collection has more than one hundred thousand posters from around the world and continues to grow. Thirty-three thousand are available online. Posters from the United States, the United Kingdom, Germany, Russia/Soviet Union, and France predominate, though posters from more than eighty countries are included.

Artificial Intelligence and the Future of War

By Bradley L. Boyd

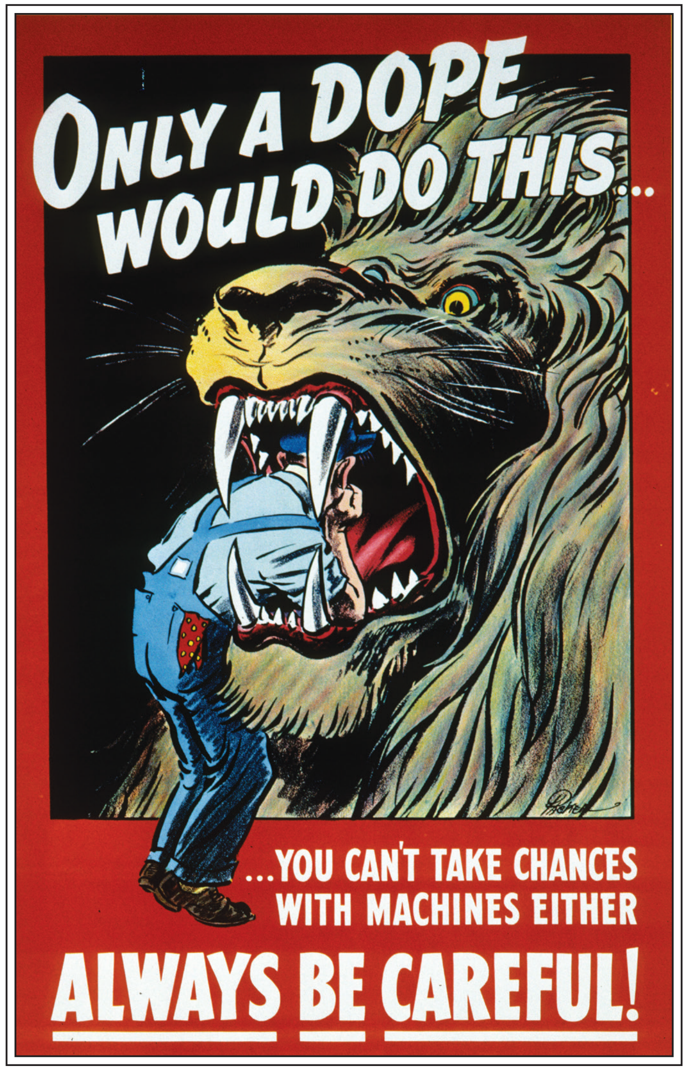

Image credit: Poster Collection, 06817, Hoover Institution Archives.

The wars in Ukraine and Gaza have shown the world that a new technology is now affecting the battlefield: artificial intelligence (AI). The Ukrainian and Russian armies are using AI to help locate and identify targets, pilot drones, and support tactical decision-making. Similarly, the Israel Defense Forces have used AI in those roles to identify and locate Hamas leadership, command centers, and patterns. The introduction of AI to warfare has shown that the technology can and will increase the speed and scale of violence, resulting in effects that are both impressive and troubling. Integrating AI into warfare renders human oversight difficult in some cases, impossible in others, and in a sense, places the conduct of war into the hands of machines, intelligent or not. No longer confined to automating manual labor, machines are beginning to automate human reasoning. Understanding how this transition will, and should, fit into war is one of the defining problems of our time. However, introducing "thinking" machines into war is not that new. Militaries have been attempting it since World War II at least.

One of the most famous founding fathers of computer science, Alan Turing, became a household name by imagining how machines might think for humans and thus solve problems at greater scale and speed than humans can. The first pieces of this work were at Bletchley Park where Turing and his companions were working on breaking the German Enigma Code using new computers like the Bombe and the Colossus. Turing's following work on computers as thinking machines made his name synonymous with the test for determining whether a machine had achieved consciousness or not. After Turing, both the British and the Americans continued work on integrating thinking machines into warfare. The U.S. Navy created a landmark system in 1958 called the Perceptron that was intended to learn tasks like a human. *The New York Times* reported that Perceptron was expected to learn to "walk, talk, see, write, reproduce itself, and be conscious of its existence," in short order. As with most of the history of AI, the hype proved overly optimistic. Perceptron and other systems failed to deliver thinking machines in the 1950s and 1960s, sending the field of AI into its first "winter."

However, the concurrent revolution provided by innovations in microchips still led to the introduction of microprocessors into military weaponry. This resulted in precision guided munitions that could hit targets on their own once fired but also in systems that could fire automatically. Many of those systems are still with us, like the Navy's Close-In Weapon System (CIWS) and smart anti-ship mines, and the Army's Patriot Air Defense System. All can engage targets on their own without further human intervention once activated if allowed to do so, meeting the U.S. military definition of autonomy. None of these systems has any AI in them, in a modern sense, and no one would accuse CIWS or Patriot of "thinking." Further, none of these systems has caused global concern over the last few decades even when they are responsible for catastrophic mistakes. Something in our perception of the technology has changed. Capabilities we now refer to as AI in military systems seem much more promising in their effects, and much more troubling.

Modern AI technology promises powerful improvements to any military fighting wars. For most of the history of war, technology has done some level of automating human labor. Whether throwing things, digging things, building things, or smashing things, machines were mostly used for labor. With the advent of computers, the idea of automating human cognitive labor seemed within reach. Whether PowerPoint is in fact a labor-saving device remains controversial but certainly transmission and storing of military plans and products with computers was helpful, and the value of precision munitions is without dispute. In the last ten years, the introduction of AI into software seemed to promise something new: the automation of reasoning.

Over the last several decades of war, the availability and sheer amount of warfighting information increased the cognitive load on military personnel beyond their ability to accomplish all that could be accomplished. As a result, the U.S. military's first AI project in 2017, Project Maven, formed to help intelligence analysts sift through all the imagery and video gathered by thousands of sensors. This seemingly simple project, identify objects in digital files, was poised to transform the speed, precision, and resilience of military operations across the force by reducing cognitive load through the automation of low-level human reasoning. As Project Maven churned through this effort, new technologies began to appear that promised even more.

The generative AI revolution in 2022 supercharged discussions of what AI could do. Suddenly, simple machine reasoning, or what Andrew Ng described as automating what would take a human about one second to process, became a low bar even though no one is able to jump over it consistently. New Large Language Models (LLMs) seem to perform very advanced, human-like reasoning and suddenly the idea of using machines to help humans make better military decisions, or even make military decisions on their own, appears within reach. These assumptions led to the development of several types of systems that are currently under development in most advanced militaries.

Most AI-enabled military systems use one of a few distinct but related technologies. First, object classification, sometimes colloquially referred to as computer vision, is used to identify things. This is the basis for the military capability called Automatic Target Recognition (ATR). ATR is the wellspring from which almost every AI-enabled system flows. Most military problems begin with being able to find and confirm a target faster and more reliably than your enemy. From drones like the Switchblade to targeting software like Gospel, an AI model is being used to find and identify the needle in the haystack of data.

Second, insight generation from large data sets is where AI models identify patterns that humans cannot see because the size of the data set means it would take too long for a human to make the same connection, if ever. Insight generation is where LLMs get their real mojo from as well, but what makes them so powerful is that the interface is in natural language. You don't need a spreadsheet; you just ask a question, and the model generates the insight. There is some overlap here but at its simplest any LLM is about making it easier for humans to access complex models without expertise to gain those insights and pattern recognition. This feels like interacting with an intelligence, but it is not.

This type of technology is very prominent in decision support systems (DSS) and, coupled with ATR, can make human decision-making faster and possibly higher quality. The ultimate prize in an AI-enabled DSS is to achieve the Move 37 effect, which refers to the AI Model AlphaGo that defeated the Go player Lee Sodol in 2016 by using a game strategy never before seen by humans and therefore unanticipated and difficult to counter. This is catnip for military leaders and the ultimate decision advantage in war.

Third, learning behavior is where an AI model learns how to behave in ways that help complete a task. Object avoidance and target tracking in drones are good examples. The average human cannot hope to fly a drone through a forest at a target but with object avoidance and target tracking, the AI helps smooth out your mistakes so you can focus on your target and not become distracted by avoiding trees.

Fourth is a combination of many technologies resulting in some type of mission autonomy. An autonomous system would likely never make use of one AI technology. Instead, it would be a few AI technologies like object avoidance and ATR in a drone that allow the system to act without human supervision. While difficult to define, autonomy in some forms is making its appearance on the battlefield.

AI-enabled autonomy can be described in three forms. Closed-loop systems are those that are a single weapons system that performs its own functions. A lethal autonomous drone is an example. It navigates itself, finds the target, computes a way to engage the target, decides to engage, shoots, then assesses its actions. This type of system is desirable because it allows for speed and resilience of mission, but it is very hard to get right.

A more likely type of AI-enabled autonomy is the distributed system. These systems are a patchwork of many systems. Perhaps a remotely piloted drone has ATR on it that can send target info to an AI-enabled DSS in a command center that determines the best allocation of resources for engaging thousands of targets and, after deciding which weapon is the best shooter, sends the firing order to a missile battery that uses AI to decide which is the safest position to fire from, providing its human operators with direction and guidance on where to move and when to fire.

A distributed system has humans involved, and humans can be the chokepoint preventing decisive speed to defeat the enemy. One way to remove the slow humans is through agentic warfare. An AI agent is a DSS that can stack complex series of tasks to perform several missions in support of a higher human intent. In the above scenarios the AI-enabled DSS could be an AI agent that has been told to establish air superiority and then constructs a task-stack to achieve that goal, controlling many systems along the way.

If this all sounds unsettling, good. It is. There are some problems with AI technology that prevent these systems from being the flawless partners in war that we might want. AI technology, of any currently fielded variety, is a probabilistic tool, meaning everything it outputs is based on a probability of being the desired output. Traditional software is deterministic, which means that if you enter the same input 100 times, you will get the same output 100 times. Not so with probabilistic software like AI.

AI learns from a set of training data how to give you the output you want. This is a strength in that it can learn complex things like how to pick out a dog from a cat. But it has a problem called an asymptote. The model will never be 100 percent certain that its output is correct. And the more the input varies from the training data, the more likely the output will be incorrect.

War is incredibly complex and there are many instances where real world input is far from the AI's training data. This renders the model unreliable or makes it reliable in very narrow circumstances. Having to employ an AI in narrow contexts or face severe reliability issues means that humans must retain control to an extent that renders the machine more a hindrance than an asset. This is why drones in Ukraine remain piloted by humans rather than AI.

The potential of AI-enabled systems in war is so promising though that no military will forgo developing these systems. So, what is to be done? The U.S. military has initiated a policy that lethal autonomous systems, powered by AI, must reflect appropriate levels of human judgment. The international community insists that any AI-enabled weapon must be under meaningful human control. But what is meaningful control, and when does meaningful control sacrifice meaningful autonomy to the point where having the system is no longer useful? When must judgment or control be applied? Continuously? Only when conducting a strike? When the system is being manufactured? There are no clear answers to these questions and it is likely that militaries will be forced to make hard choices based on the severity of their security situation. In the meantime, the machines are learning, and war is growing faster and ever more frightening.

Bradley L. Boyd was a visiting fellow at the Hoover Institution. Boyd's research interests focus on the integration of emerging technology into warfare and national security. He is particularly interested in the rise of automation and autonomy in military systems, to include decision making, planning, information operations, enterprise operations, and weaponry. Boyd also looks at the way U.S., Chinese, and Russian military integration of emerging technology affects military, economic, and social stability. Prior to joining Hoover, Boyd served as a defense and foreign policy advisor to Senator Angus King and as Senator King's representative to the Cyberspace Solarium Commission. Boyd also served as the director of AI-enabled warfighting capability development at the U.S. Department of Defense's Joint Artificial Intelligence Center and as the director of General Mark A. Milley's coordination group. Boyd was a senior military fellow at Stanford's Center for International Security and Cooperation and a Seminar XXI Fellow at the Massachusetts Institute of Technology. Boyd is a former U.S. Army and U.S. Marines officer with operational experience throughout the Indo-Pacific, Europe, Middle East, and Central Asia. Boyd has a master's degree in international politics from the University of Cambridge, a BA in anthropology from the University of California–Irvine, and a master's in cyber strategy from the U.S. Army's Command and General Staff College.

Regulating New Weapons: Lessons for the Age of AI

By Lena Trabucco

Image credit: Poster Collection, 02189, Hoover Institution Archives.

Introduction: Weapons, Ethics, and the Burden of Law

Throughout history, the creation of new weapons has unsettled the assumptions of statesmen, commanders, and jurists alike. The crossbow was once denounced as inhumane, aircraft were considered too dangerous, nuclear weapons forced the creation of entirely new doctrines of deterrence, and chemical agents demanded the strengthening of international treaty law. At every turn, the legal frameworks governing war—the Hague Conventions, the Geneva Conventions, and the customary law of armed conflict—have struggled to adapt to technologies that hold particular military significance. Artificial intelligence (AI) now presents a comparable challenge. Autonomous systems capable of selecting and engaging targets raise questions about the long-standing rules of accountability and the capacity of AI to comply with core international humanitarian law (IHL) obligations, such as proportionality and distinction.

Central to the concern of AI-enabled autonomous weapons (AWS) is the issue of human control. This refers to the degree of human judgment required to guide and, potentially, contain the performance of an AWS in a combat environment. But there are many ways to exhibit human control that go beyond the role of operators who are tasked with supervising a deployed AWS. Legal advisors and legal actors, more generally, are an important element of human control with military AI. They are tasked with evaluating the performance of an AWS against the legal obligations of a particular nation, to ensure the new weapon system can function in compliance with IHL. When it comes to AWS, this is a challenging task.

The legal community has faced similar dilemmas before, and the mixed record of regulating expanding bullets and chemical weapons provides valuable lessons for AI regulation. Importantly, this does not mean AI weapons are technically comparable to expanding bullets or chemical weapons—far from it. Nonetheless, it is useful to examine the experience of regulation that evolved from these new capabilities to identify lessons for regulating AI-enabled warfare.

Expanding Bullets: Early Humanitarian Law in Practice
A precedent for prohibition, but fragile without enforcement

The late nineteenth century brought a debate over “expanding” bullets (also called “dum-dum” bullets), which expanded upon impact, causing horrific wounds. In 1899, the Hague Peace Conference prohibited their use on humanitarian grounds. This was one of the first times states codified a restriction not for military necessity but for the principle of limiting unnecessary suffering, a significant milestone in the evolution

Poll Question: Does the use of artificial intelligence in modern weaponry constitute a revolution in military affairs?

- ☐ Artificial intelligence (AI) has transformed modern warfare and threatens to make weapons autonomous without the oversight of the operator.
- ☐ AI will force nations to recalibrate their defense budgets and usher in fundamental changes to weapon systems.
- ☐ AI will increase the lethality of weapons on the battlefield while simultaneously reducing civilian casualties.
- ☐ Human control over AI-enabled weapon systems will always be required either due to ethical standards or international law.
- ☐ AI is little more than the next iteration of a process of technological advancement dating back thousands of years.

of IHL. However, practice did not always match principle. Many colonial powers argued that the prohibition on expanding bullets only applied to conflicts among European nations, while loopholes and weak enforcement mechanisms undermined compliance. The takeaway from the expanding bullets experience: the mere existence of a legal instrument is not enough. For new weapons, the law requires not only codification but also credible means of verification and application across all contexts of war.

Chemical Weapons: The Power of Legal Taboo
From partial ban to comprehensive regime

World War I seared chemical weapons into the collective memory. The gas clouds over Ypres revealed the inadequacy of existing law and motivated the 1925 Geneva Protocol, which prohibited their use but left certain issues, such as stockpiling and transfer, unchanged. Many states signed, but with hefty reservations. It was later, with the 1993 Chemical Weapons Convention, that international law achieved greater reach, banning production, possession, and transfer, and empowering an international organization—the Organisation for the Prohibition of Chemical Weapons (OPCW)—with intrusive verification powers. In this case, the law evolved into something stronger—not just a formal treaty, but a taboo reinforced by monitoring institutions and near-universal condemnation. This trajectory demonstrates that effective regulation of new weapons depends not only on legal rules but also on the institutional capacity to enforce them and the normative power to delegitimize their use.

The Legal Obstacles to Regulation
Ambiguity, dual-use, strategic rivalry, and verification all undermine law

The difficulty of regulating AI under international law lies in four interlocking problems. First is definitional ambiguity. Unlike expanding bullets or chemical agents, there is no universally accepted definition of what qualifies as an autonomous weapon system, and States continue to be divided. Second is the problem of dual use—algorithms developed for benign civilian purposes can be repurposed for war, making categorical bans difficult. Third, strategic competition discourages restraint. Major powers fear that restrictions will leave them vulnerable in an AI arms race. Finally, there is the problem of verification. Whereas chemical stockpiles could be counted and destroyed under OPCW supervision, algorithms leave few physical traces and often lack transparency.

There is also the issue of mission-critical technologies. Despite expanding bullets and chemical weapons seeing a certain degree of success in legal regulation, many other technologies throughout history have not moved the legal needle toward enhanced regulation due to the mission-critical nature of the weapon.

Regardless of efforts toward prohibition or restraint, when a technology becomes central to military power, efforts to ban it usually collapse.

Conclusion: Law Must Anticipate, Not Follow, Catastrophe

The lesson of history is clear: waiting for disaster leaves law behind

Artificial intelligence presents a complex challenge for international law. Unlike chemical agents, which could be clearly identified and banned, AI is not a single substance or weapon but a suite of dual-use technologies, many of them developed in civilian industries.

The regulation of expanding bullets shows how fragile humanitarian law can be when enforcement is weak. The Chemical Weapons Convention demonstrates how law can be strengthened through robust institutions and the establishment of international taboos. AI weapons fall somewhere in between: too diffuse to ban outright, too consequential to ignore. The international community thus faces a choice: to adapt the law of armed conflict proactively, find a middle ground through policy for regulation, or wait for the first catastrophe involving AI weapons to trigger political will.

Lena Trabucco is an assistant professor at the Centre for Military Studies at the University of Copenhagen. She is also a non-residential research fellow at the Stockton Center for International Law at the U.S. Naval War College and a research fellow with the Tech, Law, and Security program at American University.

Image credit: Poster Collection, 04243, Hoover Institution Archives.

Human "Control" over Autonomous Weapons Systems

By Ariel Conn

Artificial intelligence (AI) is already used in a range of military applications, from logistics to navigation to HR, but the most debated and the most legally and ethically contentious application is, of course, the use of AI in weapons systems. A primary source of concern, hope, and hype around autonomous and AI-enabled weapons systems (AWS) is the extent to which they can function outside of traditional human control.

For over a decade, an international debate has transpired around what human control over AWS must look like. Many argue that weapons that can function autonomously, especially with respect to selecting and targeting humans, are inherently illegal and unethical: only humans can be held legally responsible and accountable, and only humans can make ethical decisions.

Yet, as these legal debates have transpired, militaries and weapons manufacturers have rapidly added increasingly autonomous features to existing weapons systems, while also developing new, AI-enabled weapons. Given the international consensus that only humans can be held responsible and accountable for weapons systems, the question of what human control or oversight looks like becomes ever more pressing.

Even if the system does what it's supposed to, there are concerns about human control and human authority over the actions. However, the question becomes especially relevant when considering what might happen if an AWS launches an attack against the wrong target, against civilians, or against people or infrastructure that are protected by International Humanitarian Law (IHL).

The Challenge of Human Control vs. Autonomy

Defining and ensuring human control over an AWS is especially tricky because autonomy exists on a spectrum with total human control on one side and full autonomy on the other.

For example, a knife is fully under human control, whereas, with a gun, once the trigger is pulled, humans no longer have control over the bullet. Drones, meanwhile, may be a mix of algorithms and remote piloting, or they may mostly fly and function autonomously, and airplanes can practically fly themselves. The IAI Harpy has been around for many years and, like other loitering munitions, is basically an autonomous system, though technically, its sensors cannot target humans directly. An increasing number of systems exist now that can function autonomously and could target humans.

Many people discuss AWS as if it's a new category of weapons systems, however, AWS is actually any weapons system built with any number of autonomous and AI-enabled capabilities. As a result, rather than representing an easily defined, new category of weapon, AWS can include a huge range of weaponry, sitting across most of the controlled-to-autonomous spectrum.

Examples of technologies that can increase autonomous capabilities in AWS include: sensors for obstacle avoidance, GPS navigation, sensors for target identification, algorithms for target identification and facial recognition, algorithms to help systems fly smoothly even when remote controlled, and communication technology to relay information to humans or for system-system communication, such as swarm technology.

With the adoption of such autonomous capabilities, weapons can be more readily deployed in locations that humans can't access and in numbers far greater than the number of soldiers on the ground. How can human soldiers and commanders possibly maintain control of systems that are deployed outside of communication or in numbers too great for human oversight?

One answer is to consider the humans involved beyond just the time of use.

The Role of Humans Across the Weapon's Lifecycle

Conventional weapons systems have taught us to look at the time of use. When something goes wrong with a conventional weapon, the problem is likely due to user error, a faulty system, or malicious intent. However, if something goes wrong with an AWS, because it can function so far outside of a commander's or soldier's control, determining who is responsible by looking only at the time of use will too often prove to be a futile task.

Instead, with AWS, it is more critical than ever to look at the full lifecycle of the weapon system.

For an autonomous system, the time of use is when the system is in its most autonomous state. However, by looking back across the timeline of development, human control increases. Using the IEEE-SA Lifecycle Framework as an example, humans have the most control during political decision-making, ideation stages, and early research and development stages. As the system is developed, human control shifts to responsibilities around testing and assurance of the system.

Human control and responsibility also include providing human users and commanders the training and consideration necessary to prepare them to use the system correctly and to recognize signs that something is going wrong with the system. Processes like Human Readiness Levels offer one method for ensuring that these practices are built into the design and development stages of the system and that the AWS does not move forward in its lifecycle until these various assurances have been met.

This approach is not without precedent: a number of lessons can be learned from the highly automated and autonomous aerospace industry. For the purpose of this essay, the most important lesson is that, when a plane crashes, the response is consistently to determine: 1) if the problem was due to user error or technical error, and 2) to look back at logs to identify who was responsible for either insufficient training for the user or for the technical components that caused the malfunction.

Human control over AWS is an understandably contentious topic, and in order to meet legal requirements and ethical guidelines, some level of control must be maintained. This may be nearly impossible to achieve when looking only at the time of use, but it may be possible to define and establish human control—and more importantly, human responsibility and accountability—when the full lifecycle is taken into account.

Ariel Conn is an instructor in computing ethics at the University of Colorado Boulder and the founder and president of Mag10 Consulting, where she works with researchers and not-for-profit organizations to address a variety of global ethical and policy issues. Her work is primarily focused on the use of AI in the military and other ethically challenging uses and impacts of AI. She is an expert advisory member for the Global Commission on Responsible AI in the Military Domain, and she was one of 2023's 100 Brilliant Women in AI Ethics.

Discussion Questions

1. Does the use of artificial intelligence in modern weaponry constitute a revolution in military affairs and threaten global security?
2. Should the use of artificial intelligence in modern weaponry be regulated by international law and monitored by international organizations?
3. Do the United States, Russia, and China utilize artificial intelligence differently in the development and deployment of weapon systems?
4. Is the United States leading the way in the development and use of artificial intelligence in modern weaponry, and if not, which nation or nations pose the greatest threat?
5. What ethical issues arise surrounding the use of artificial intelligence in weapons, and how important is human oversight of those weapons?

Military History in Contemporary Conflict

As the very name of Hoover Institution attests, military history lies at the very core of our dedication to the study of "War, Revolution, and Peace." Indeed, the precise mission statement of the Hoover Institution includes the following promise: "The overall mission of this Institution is, from its records, to recall the voice of experience against the making of war, and by the study of these records and their publication, to recall man's endeavors to make and preserve peace, and to sustain for America the safeguards of the American way of life." From its origins as a library and archive, the Hoover Institution has evolved into one of the foremost research centers in the world for policy formation and pragmatic analysis. It is with this tradition in mind, that the "Working Group on the Role of Military History in Contemporary Conflict" has set its agenda—reaffirming the Hoover Institution's dedication to historical research in light of contemporary challenges, and in particular, reinvigorating the national study of military history as an asset to foster and enhance our national security. By bringing together a diverse group of distinguished military historians, security analysts, and military veterans and practitioners, the working group seeks to examine the conflicts of the past as critical lessons for the present.

Working Group on the Role of Military History in Contemporary Conflict

The Working Group on the Role of Military History in Contemporary Conflict examines how knowledge of past military operations can influence contemporary public policy decisions concerning current conflicts. The careful study of military history offers a way of analyzing modern war and peace that is often underappreciated in this age of technological determinism. Yet the result leads to a more in-depth and dispassionate understanding of contemporary wars, one that explains how particular military successes and failures of the past can be often germane, sometimes misunderstood, or occasionally irrelevant in the context of the present.

Strategika

Strategika is a journal that analyzes ongoing issues of national security in light of conflicts of the past—the efforts of the Military History Working Group of historians, analysts, and military personnel focusing on military history and contemporary conflict. Our board of scholars shares no ideological consensus other than a general acknowledgment that human nature is largely unchanging. Consequently, the study of past wars can offer us tragic guidance about present conflicts—a preferable approach to the more popular therapeutic assumption that contemporary efforts to ensure the perfectibility of mankind eventually will lead to eternal peace. New technologies, methodologies, and protocols come and go; the larger tactical and strategic assumptions that guide them remain mostly the same—a fact discernable only through the study of history.

Hoover Institution, Stanford University
434 Galvez Mall
Stanford, CA 94305-6003
650-723-1754

Hoover Institution in Washington
1399 New York Avenue NW, Suite 500
Washington, DC 20005
202-760-3200

HOOVER INSTITUTION