MODERN WAR INSTITUTE AT WEST POINT

THE MWI AUTONOMY PAPERS

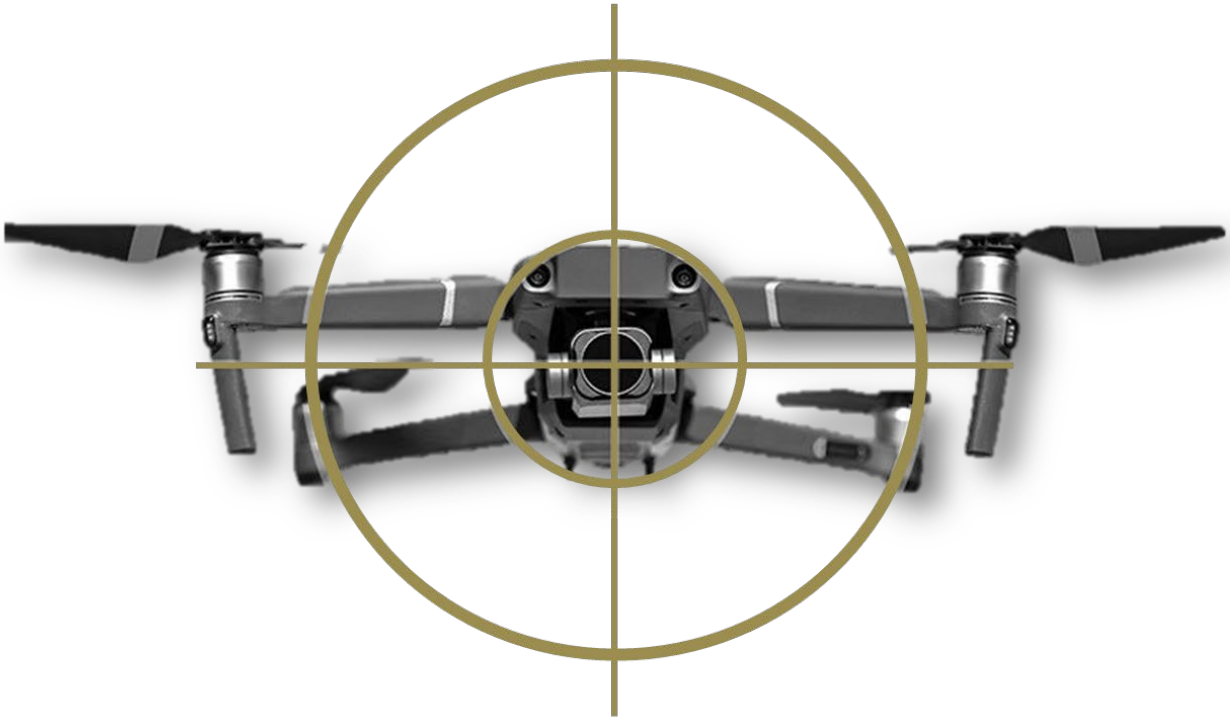

Recent conflicts have provided ample examples of how autonomous systems are changing warfare. During 2025, MWI focused its research efforts on understanding how autonomous systems will contribute to a future fight. MWI conducted a contemporary battlefield assessment of the Russo-Ukrainian War to identify lessons. These efforts culminated with the annual Modern War Conference on "Autonomy in the Future Fight."

To support that conference and share the work of our many contributing authors, we have collected a selection of articles on artificial intelligence and unmanned systems as well as how to counter those systems. We will publish additional articles in the winter 2026 edition of the Modern War Journal.

The point of contact for this collection is Major Robert G. Rose, the Modern War Institute's chief research officer, at robert.rose@westpoint.edu. MWI published these articles using hyperlinked citations that do not appear in print. You can find the citations on the MWI website at https://mwi.westpoint.edu/the-autonomy-papers-2025/ or using this QR code:

Contents

Artificial Intelligence

Unmanned Systems

Countering Unmanned Systems

Disrupting Systems: Cyber and Electronic Warfare

Artificial Intelligence

Tools Are Not Strategies: A Short Guide on Artificial Intelligence for Defense Professionals

Jules Hurst

Humans have a tendency to treat emerging technologies as catch-all solutions to endemic problems. Information technology companies actively reinforce this notion in their marketing. Google, Amazon, IBM, and other firms do not just offer software and hardware, they offer "solutions." Unfortunately, technologies are never solutions in isolation. Technologies are tools and these tools (should) form part of an organization's strategic approach to solving enduring problems. They are not solutions on their own, and they are certainly not strategies.

Selecting the wrong tool—or even the right tool's misapplication—does more harm than just squandering resources. The use of one technological tool prevents use of another because of the time and money it takes to incorporate it into workflows, train staff on its use and maintenance, and acquire equipment or services that support it. Opportunity costs are substantial.

As defense leaders look toward artificial intelligence and machine learning (AI/ML) to address organizational problems, they must know more than the hype. They must be familiar with the strengths and weaknesses of AI/ML techniques, consider other tools and steps to solve their issues, and understand the resources required to implement AI/ML tools into approaches. AI/ML algorithms can have great effects on organizational processes, but only if their strengths align against issues within those processes. Even then, AI/ML algorithms may not be the best tool for the task.

Establishing a starting point for military professionals to begin thinking through these requirements is imperative, and there are several rules of thumb that can help leaders gauge whether their organizations should use AI/ML as a tool in their organizational approaches.

Pattern Matching

Current AI/ML techniques create algorithms with the ability to detect patterns. This seems trite, but pattern matching is fundamental to intelligence. From birth, each of us begins the development of our own pattern-detection skills that lead to our ability to identify objects, make predictions, draw conclusions, learn language, and interact with our environment. With time, experience, and exposure to exhaustive datasets, our minds construct the heuristics and algorithms that some of the world's brightest minds want to replicate through computation.

As babies and toddlers, humans gain the ability to identify everyday things by discerning common characteristics in classes of objects. Our ability to make greater distinctions rises with exposure to more examples (data). To a one-year-old, every truck or van is a car. By three years old, children may recognize individual brands, colors, and models. Parents can supervise this learning by intentionally identifying examples of cars or allow the learning to occur unsupervised as children watch others identify and interact with automobiles.

With appropriate amounts of data, machine-learning algorithms can replicate this pattern recognition. Like humans, algorithms "learn" to recognize more detailed characteristics when trained on more data (but also like humans, not indefinitely). Algorithms can exceed human capacities to identify objects or patterns at scale, but not necessarily in accuracy. For problem sets that require the triage of masses of data (x-rays, magnetic resonance imagery, satellite images, radar tracks, sentiment, etc.), machine-learning algorithms can help if definable signatures exist. AI/ML models trained to detect tumor growth are already showing promising results. Most interestingly, algorithms can run in congress with one another to perform more complex tasks in ways that simulate general intelligence.

Machine-learning algorithms only reliably identify patterns when used in operational environments consistent with their training environment. Today's algorithms cannot dependably generalize. An algorithm must experience, while being trained, whatever variances users expect it to encounter in operation. An algorithm trained to identify human faces in profile against a black background is just that. Something as simple as changing the background color could harm the algorithm's effectiveness. Any deviation in operational conditions from training conditions impacts algorithmic accuracy. A 2019 paper highlighted the failure of a Google algorithm to reliably identify objects outside of normal orientations; the algorithm could easily identify vehicles driving on the road, but as soon as the vehicles flipped on their axis, the algorithm made startling errors. Humans can generalize after a few examples. Algorithms cannot. They need explicit instruction and training, at least until things like zero-shot learning mature.

Algorithmic inability to generalize explains why algorithms can successfully identify tumors on MRIs—because the data used to train the algorithm perfectly matches the operational data—but struggle in complex, live environments. Machine-learning algorithms are wonderful pattern matchers, but only when conditions for operational use match well with training conditions. If a problem set exists in a tightly controlled environment and training data contains the same variances the algorithm will experience operationally, it will likely work well. AI/ML algorithms trained to play games outperform human competitors for this very reason—the conditions of a game or simulator are controlled; the operational and training environment are the exact same. If an algorithm's operational environment contains more variance than its training data, grievous errors will occur.

The accuracy of algorithms in pattern recognition, to include classifiers that identify objects, do not necessarily make them good or bad. For some applications, an algorithm with 70 percent accuracy would still add value and reduce human labor. In other applications, it would not. Leaders must understand the level of accuracy necessary to add value and identify levels that would be hazardous. Individuals who work with the algorithm must know the approximate accuracy of the algorithm in various scenarios to best use it. This allows operators to avoid abdicating judgment in situations where the algorithm is prone to fail.

Acting on Information

When an algorithm can identify objects in an environment, it can form rules to interact with them, much as humans do. Once children can identify animals like bees, fire ants, or wasps, they form rules to avoid them for safety. Once machine-learning algorithms can identify objects, engineers can either prescribe rules for interacting with those objects or allow other algorithms to develop rules for interaction with that object type. A self-driving car might have a rules-based algorithm to tell it to stop when an attached sensor detects a red light (likely through machine learning), or employ an algorithm trained with the goal of avoiding car-on-car collisions.

Algorithms that identify objects should not run unsupervised unless the probable harm caused by a mistake is small relative to the utility gained, the harm is reversible, or where an operator can minimize risk through control measures. Spam filters are a good example. If an email's spam filter misidentifies an email from a friend as junk mail and places it in a spam folder automatically, the harm is small and reversible, and the user can make rules that prevent it from occurring again (control measures). From a utilitarian perspective, the good supplied by a spam filter generally outweighs the loss of an occasional email.

Alternatively, if the US Navy employed naval mines that used an algorithm to discern combatant from noncombatant vessels, the consequence of an algorithmic mistake would be much greater. Many nations would only use this type of weapon if its error rate was exceedingly small. Unfortunately, predicting the error rate of an identification algorithm in all circumstances taxes human prescience, even with extensive testing. Algorithmic reasoning is often inconceivable to human beings. Neural networks mimic human neurology, but they are not artificial duplicates. Paradoxically, thinking machines do not lessen the need for human judgment, but increase it.

Furthermore, algorithms that make a mistake lack the general intelligence, perception, or context to stop making it. A human driving a car would immediately stop if he or she struck a pedestrian. Audible signals, and the collision itself would tell the driver something has gone wrong. A self-driving car might keep driving without human intervention. A machine-learning algorithm will continue selling or buying equities without market context. AI/ML algorithms lack the general context to detect mistakes and change their behavior without explicit training.

When You Need Machine Learning—and When You Do Not

Machine-learning algorithms add the most benefit when used to conduct pattern analysis at a speed or scale that manpower constraints make prohibitive, in situations where analysis must occur continuously, or where the number of variables exceeds human capacity to analyze. Analyzing satellite imagery using machine learning is an excellent application of the technology. Government and commercial satellites generate torrents of images that require analysis by professionals at all hours. Humans can analyze images, but an appropriately trained algorithm can do this at a speed that reduces organizational manpower costs and frees human analysts for higher-level tasks. Moreover, use of an AI/ML-powered image classifier on satellite imagery can transform an imagery database into something searchable by text by turning pixels into objects; "Find all T-72 tanks identified on imagery located within 20 miles of City X between 1/2/2001 and 1/2/2002." Employment of machine learning in this context exploits its advantage in performing analysis at speed and scale and adds secondary advantages that can improve organizational performance or processes.

Frequently, leaders want to incorporate machine learning when rules-based algorithms will accomplish the same task. Recently, a senior Army officer proposed that the US Department of Defense create a machine-learning algorithm to identify individual soldiers capable of mobilizing to support civil authorities during natural disasters. A machine-learning algorithm could perform this task, but using one is more laborious than programming a rules-based query. The Department of Defense knows exactly what kind of soldier it needs to mobilize. He or she should be physically ready (which the military quantifies with metrics), be trained in a codified military skill set (which each service tracks with a system of codes reflecting occupational specialties and additional skills), be outside of the immediate disaster area, and have no administrative bars to mobilization (e.g., injured, attending military education courses, etc.). Pulling this information is no more complicated than a database query: "Find all soldiers with [specialty code or skill identifier], AND who are medically ready, AND have no administrative restrictions on mobilization, AND live outside of a 50-mile radius of a disaster area." There is no need for machine learning to predict or identify what kind of soldiers would best meet the needs of the mission. The Department of Defense has already done all the thinking and it has structured data to

query for an answer. Leaders should only reach for AI/ML models if other techniques cannot meet needs as quickly, accurately, or cost-effectively.

If a competent data scientist can solve your problem, you likely do not need to employ an AI/ML model. As you assemble the data, transform it, and analyze it to create a machine-learning model, you will arrive at an answer before you set AI to work for you. Generally, you need an AI/ML model when the quantity of data overwhelms the manpower available to accomplish the task, you require a model to continuously evaluate data for changes (e.g., credit card fraud detection), or the number of variables exceeds human capacity to analyze (e.g., games, simulations, genome decoding).

What Does It Take to Make a Model?

While AI/ML models can take large amounts of computational power to train, create, and run, the most important logistical requirement to creating an AI/ML model is the acquisition of data that matches the conditions where the model will operate. This is more difficult than it sounds.

Imagine a lawn mower engine manufacturer wants to create a predictive maintenance algorithm. The company might find maintenance records at repair shops around the globe and aggregate them to see if an algorithm can pick out patterns of mechanical failure. Repair records will exist in different formats and with different data fields, free-text problem descriptions, different languages, and varying accuracies. The engine manufacturer would need data-literate engineers to comb through the reports and fit free-text descriptions of mechanical failures into a formal ontology of mechanical problems, an arduous task.

Even if the company finds sufficient records and standardizes them, the model might be less insightful than a competent statistician. Good data hygiene and data-science techniques will often provide the same answer as machine learning. For the algorithm to add value, the model needs access to data inside engines as they fail in real time so that users can predict when parts inside individual engines need replacement. The company would need to redesign its product with gauges and sensors to supply the necessary data. Even then, the algorithm would require periodic updates. Algorithms, like people, always require training when environmental changes occur and engineers must design or redesign hardware with AI in mind.

Finding data for an algorithm to train on that closely simulates operating conditions is the hardest task in algorithm creation. An image classifier needs thousands of hand-labeled images to develop an initial capability. These images should be at the same resolution, from the same perspective(s), and within the same spectrum as images in the operational environment. The number of hand-labeled images required will likely grow with the complexity of the classification task.

As you consider use of an AI/ML tool, ask yourself the following: Do I have access to appropriate data to train the algorithm? Is this data the same as data the algorithm will evaluate in the operational environment? How many man-hours will I require to acquire, label, and format the data? Is there a simpler tool than AI/ML that achieves the same effect?

Heuristics for Heuristics

Like many new technologies, the hype behind AI/ML threatens to interfere with its objective, fact-centered implementation. Social media posts and business presentations barrage senior leaders with buzzwords that imply machine learning is a panacea. This marketing is hard to overcome. As leaders consider incorporating AI/ML technologies into their strategies for solving hard problems, let the following rules of thumb act as a guide.

- Don't be the First: If another organization has not already done something like what you would like to do with AI/ML, then let someone else work through the problem. If you are not a research and development organization, let a federally funded research and development center, university affiliated research center, or Department of Defense lab prove the concept.
- No Data, No Algorithm: If you cannot find ample data to feed algorithms that matches that of the algorithm's operational environment, your algorithm will struggle.
- Know Thyself: The algorithmic accuracy necessary to gain efficiencies varies with organizational need. Fraud detection algorithms can add value to operations with a low accuracy rate. Self-driving vehicle algorithms cannot because of the high cost of an error. You must know how accurate or precise an algorithm needs to be to help your organization. This allows you to evaluate the costs and benefits of its creation and implementation.
- Don't Overcomplicate it: Never use a more complicated tool when a simpler one will render the same result for less effort over the life of a program. Many organizational issues do not require AI/ML for resolution. Traditional, rules-based software can provide just as much of a productivity increase as machine learning tools in some situations.
- Narrow and at Scale: At present, AI/ML algorithms best perform narrow tasks in controlled or constrained environments. As the number and complexity of scenarios an algorithm may experience grows, performance declines. Algorithms provide the largest productivity gains when asked to match patterns in a narrow context at a speed, scope, or complexity that exceeds human capacity.

AI/ML technologies offer defense leaders incredible tools to attack longstanding problems, but they are not solutions in and of themselves. Prudent leaders must carefully examine the costs of developing AI/ML algorithms, their likelihood of success in implementation, and their relative advantages or disadvantages to other tools. If defense budgets decline, provident selection of AI/ML projects will become increasingly important.

Persuade, Change, and Influence with AI: Leveraging Artificial Intelligence in the Information Environment

Austin Coombs

US adversaries are weaponizing artificial intelligence to unleash a new wave of psychological warfare. Russia, through its troll factories and bot farms, has adopted a new AI-driven asymmetric warfare strategy, using generative models to amplify disinformation efforts on an unprecedented scale. A striking example was the AI-generated image of a false Pentagon explosion, which caused a rapid and dramatic (albeit temporary) drop in the US stock market. This incident highlights the catastrophic potential of AI-driven propaganda to destabilize critical systems, making it imperative for the United States to adapt. While the Department of Defense's AI Adoption Strategy is a step forward, gaps remain in training US forces to fully harness AI for information warfare and to counter these evolving threats, particularly those from Russia and China.

Russia is using AI to enhance its disinformation campaigns, particularly through the evolution of bot accounts that now produce more human-like and persuasive content. Ahead of the coming November US presidential election, Russian actors have sought to leverage AI to enhance the scope and scalability of their influence operation efforts, some of which specifically aim to shape public opinion toward candidates, sway US electoral outcomes, undermine public confidence, and sow discord both within the United States and globally. The integration of AI has allowed Russia to monitor the information environment in real time, enabling rapid adaptation of disinformation tactics.

China's use of AI in psychological warfare has become a key element of its strategy to shape regional and global narratives, amplifying its influence across the world. By leveraging AI to create deepfakes, automate social media bots, and tailor disinformation to specific audiences, China has enhanced its capacity to manipulate public discourse. This strategy extends beyond mere online influence; China's AI capabilities enable large-scale cyber-enabled operations, as seen in coordinated disinformation campaigns targeting Western audiences. China's "cognitive domain operations" merge AI with psychological and cyber warfare, aiming to deter US intervention in future conflicts or polarize American society, presenting an ever-growing challenge to global stability.

The Dangers of Doing Nothing

Failing to act against adversarial AI-enhanced information warfare poses significant risks. Russia's and China's ability to leverage AI to amplify their propaganda and disinformation campaigns threaten to undermine US and allied efforts across the globe. If unchallenged, this technological edge could enhance adversarial aims to destabilize regions, influence elections, and manipulate public opinion with unprecedented effectiveness. The cost of inaction is high, potentially leading to a strategic imbalance favoring adversaries that are adept at exploiting AI for malign purposes. Imagine a scenario where Russian and Chinese AI-driven disinformation campaigns go unchecked. The flood of false narratives could have devastating effects by eroding public trust in democratic institutions, creating confusion and division. In such an environment, the United States' ability to project influence and reinforce stability in regions across the globe could be severely diminished. The stakes are high, and the need for a proactive response is urgent.

Enhancing Training: Adapting to AI Integration

The training focus for US military personnel, particularly psychological operations soldiers, needs to adapt to the evolving technological landscape. Soldiers must be educated on the current AI tools available and understand how these tools can then assist in analyzing the operational environment, speeding up analysis, generating content, and addressing risk concerns for commanders. Increasing AI literacy is the first step. Soldiers should understand the basics of AI, its capabilities, and limitations. This foundational knowledge is crucial for effectively integrating AI tools into operations. Training programs should include hands-on experience with AI tools, allowing soldiers to practice using these technologies in realistic scenarios. Education on the ethical implications of AI use in military operations is also essential to ensure compliance with legal and moral standards. Given the rapid pace of AI development, training programs must emphasize continuous learning and adaptation to keep pace with new advancements.

AI-Enhanced Psychological Operations

Integrating AI tools into military operations, particularly in the realm of information warfare, offers several key advantages that can enhance US military capabilities and enable psychological operations soldiers to counteract adversarial information campaigns. These advantages include enhanced analysis, speed and efficiency, scalability, and risk mitigation. AI can analyze vast amounts of data from various sources to identify trends, sentiment, and potential threats. This capability allows psychological operations detachments and teams to gain a deeper understanding of the operational environment and develop more precision-based messaging efforts. Additionally, AI can generate content quickly and efficiently, which is vital in today's increasingly fast-paced information environment.

Moreover, a major issue that psychological operations teams face is the expansion of their efforts across the entirety of their deployed areas of responsibilities. Army psychological

operations capabilities are already in high demand by geographic combatant commands, theater special operations commands, and Department of State embassy country teams. Given the high demand from multiple organizations, detachments are required to break into smaller subunits to cover extensive geographic areas and critical missions, leading to substantial challenges in managing bandwidth and scalability. AI can help overcome these bandwidth and scalability issues by streamlining content production and distribution, allowing smaller teams to support wider mission objectives, cover more ground, and engage with multiple audiences without sacrificing speed or quality. This scalability is essential for countering widespread disinformation campaigns effectively when timing is usually a crucial factor in messaging effectiveness.

Risk management is also significantly enhanced by AI. AI can assess the potential impact of different messaging strategies, helping commanders to understand the risks and benefits associated with various courses of action. By simulating potential outcomes, AI can provide a clearer picture of the operational environment and the likely responses from adversaries and other audiences. Moreover, AI's risk mitigation capabilities enable teams to derive actionable insights and recommendations, streamlining planning processes to better support their commands. This is predicated on the AI implementors being able to communicate the integration process to their commanders and policymakers.

Content generation is another critical area where AI can be beneficial, as it can be used to efficiently create authentic and realistic material rapidly enough to maximize impact. AI can generate high-quality content at scale, which is crucial for countering adversarial narratives and disseminating US messaging, enabling rapid responses to adversarial propaganda. Tools like natural language processing can create persuasive and contextually relevant content that resonates with target audiences. The speed at which it can do so is crucial in the fast-paced information environment, where timely interventions can make a significant difference.

Audience response testing is another area where AI can be invaluable, primarily due to its speed and efficiency. Instead of relying solely on traditional methods, AI can simulate expected audience reactions and engagement metrics based on preloaded audience characteristics, allowing teams to refine messaging strategies before wider dissemination. While this approach may not replace the nuanced, experience-based insights of a psychological operations detachment, it significantly accelerates the process, enabling multiple iterations of a message to be tested and optimized more quickly than manual methods allow, thus improving the likelihood of effective engagement.

Training Proposal: Developing a Period of Instruction

To effectively integrate AI tools into military operations, a comprehensive training program is essential. This program should include essential instruction blocks covering the fundamentals of AI, basic knowledge of AI literacy, how AI and large language models works, various capabilities AI can provide, and, crucially, its limitations and concerns about its use. This foundational knowledge is critical for understanding how AI can be applied in military contexts. Hands-on training should be a significant component of the program. Practical exercises that allow soldiers to use AI tools in simulated scenarios will help them become familiar with the technology and understand how it can be applied in real-world operations. This hands-on approach ensures that soldiers are not just theoretically knowledgeable but practically skilled in using AI tools. Ethical and legal considerations should also be a key part of the training; soldiers must be aware of the potential risks and ensure that their use of AI complies with any strategy documents or policy updates that dictate ethical standards of AI usage.

Continuous learning is essential given the rapid pace of AI development. Ongoing education and training are crucial to ensure that soldiers remain proficient in using AI. This could include regular updates on new AI tools and technologies, as well as refresher courses to keep soldiers informed about the latest developments in AI. Specialized training for psychological operations personnel is also necessary, given their role in challenging adversarial narratives in the information environment. Focused training on how AI can enhance previously discussed specific tasks—information analysis, content generation, and audience engagement—will equip psychological operations teams with the skills they need to effectively integrate AI into their operations. The quality of this specialized training will be greatly enhanced if it can include real-world examples and case studies to illustrate the successful practical application of AI, as well as lessons learned from implementation and experimentation struggles.

Policy Updates for End-User Implementation

One of the critical solutions to countering adversarial AI advantage is updating US military policies to provide clear boundaries for the use of AI tools. Training on AI is fundamental, but its impact will only be maximized if the right policy framework is in place. Current policies often lack the specificity needed to guide military personnel in the ethical and effective use of AI technologies. By establishing comprehensive guidelines, the US military can empower its members to utilize AI in ways that support US goals and objectives while maintaining adherence to ethical standards. These policy updates should focus on defining acceptable uses of AI in various military operations, establishing protocols for the deployment and oversight of AI tools, and providing a framework for continuous evaluation and adaptation of AI policies as the technology evolves. Clear guidelines will not only enhance operational effectiveness but also ensure that AI use is responsible and ethical.

The risks of ignoring the AI-driven psychological warfare tactics employed by Russia and China are not just theoretical—they are already unfolding. As the operational environment continues to evolve, adversaries will continue to exploit AI to destabilize democratic systems, manipulate public opinion, and undermine US influence on the global stage. The cost of inaction is severe, as AI accelerates the scale and sophistication of disinformation campaigns in ways we are only beginning to grasp. Failure to address these tactics could lead to a strategic imbalance that weakens the United States, leaving us vulnerable to further erosion of trust in our

institutions and a diminished ability to project influence and reinforce stability across the globe.

The US military cannot afford to lag behind in this critical dimension of the information environment. To preserve our national security, we must adapt now. This requires not just policy updates, but a comprehensive approach that includes advanced training, strategic AI integration, and rapid deployment of AI-enhanced operations. By embracing AI as an active component of our psychological warfare capabilities, we can outpace our adversaries, address the capacity and bandwidth issues psychological operations forces face across the globe, and be better prepared to safeguard the information environment from adversarial malign influence. This is not a future challenge—it is a present-day battle, and the stakes could not be higher.

Targeting at Machine Speed: The Capabilities—and Limits—of Artificial Intelligence

Jesse R. Crifasi

The United States Army's ability to deliver precision fires and effects is fundamentally tied to its doctrinal targeting methodology: decide, detect, deliver, assess (D3A). Field Manual 3-60, *Army Targeting* prescribes the use of D3A as an integrative approach requiring cooperation across multiple warfighting functions. As the Army advances under the pressures of multidomain operations as its operational concept, optimizes its contributions to US strategic competition with near-peer adversaries, and pursues its recently announced transformation initiative, the necessity of integrating artificial intelligence into targeting workflows is paramount.

AI technologies have already proven their utility across a range of defense applications, including intelligence, surveillance, and reconnaissance processing, decision support, and autonomous systems operations. Over the past several years, a growing body of academic research has explored these capabilities, yielding insights with significant implications for military policy and doctrine. Key takeaways from this body of work include:

1. **AI in targeting presents a moral dilemma**—it must be employed as a tool, not as a substitute for the warfighter's judgment.
2. **Time is the most compelling performance metric** for evaluating AI effectiveness in the targeting process.
3. **AI offers undeniable scaling advantages**, particularly in data processing and decision acceleration.
4. **Human commanders must remain the final arbiters of lethal force**, preserving the principle of human-on-the-loop decision-making.
5. **AI should augment—not replace—critical targeting functions**, such as rules of engagement validation, proportionality assessments, and determinations of military necessity.

Even with these insights established by research, AI's integration into the D3A targeting methodology remains underdeveloped in operational doctrine. There is therefore a central question the Army has yet to answer: Can AI enable the D3A cycle to achieve faster, more reliable, and more effective targeting—while preserving accountability through human oversight?

Emerging programs such as the Israeli AI-enabled system known as "the Gospel," the US Department of Defense's Project Maven, and other kill chain automation initiatives reflect a growing desire to accelerate targeting cycles. These efforts are largely aligned with the F2T2EA (find, fix, track, target, engage, assess) model used in joint targeting. The Army, however, continues to rely on D3A as the doctrinal cornerstone of fires and effects integration at the brigade and division levels. To adapt AI to the D3A methodology, a modular and doctrinally grounded approach is required—one that maps AI capabilities to discrete phases of the targeting cycle and identifies value-added contributions to each step.

Building on a recent study in the *Naval Engineers Journal*, which mapped AI methods to F2T2EA kill chain functions, we can extrapolate AI applications for each phase of D3A. In the *decide* phase, tools such as game theory models, decision trees, and logistic regression algorithms can support enemy course-of-action development, attack asset prioritization, and effects determination. During *detect*, AI excels at target recognition via pattern association and anomaly detection, leveraging sensor fusion and deep learning. This will greatly reduce the time it takes to determine a target's functional characterization, especially at large scales. For *deliver*, optimization algorithms and prescriptive analytics can refine weapon-target pairing and target engagement timings. This has the potential to eliminate human-introduced error in the target vetting process. Finally, in the *assess* phase, battle damage estimation benefits from clustering models and explainable AI tools that support image interpretation and effects validation. Along with correlation modeling this can bring greater transparency to the combat assessment process, saving precious time and munitions as well as bringing greater clarity to the commander's decision support tools.

The study published in the *Naval Engineers Journal* examined eight specific AI methods relevant to targeting, using

empirical methods of analysis. These included logistic regression, linear regression, clustering, association rule learning, and several others. However, not all methods proved suitable for targeting. Random forest and generative adversarial networks were dismissed for their black-box characteristic—when tasked with justifying and rationalizing their targeting solutions, these systems were incapable of explanation. This is not just a technical problem but, more importantly, a legal showstopper. As those familiar with theater rules of engagement and the law of armed conflict know, commanders are directly responsible for ensuring compliance with the tenets of military necessity and proportionality when conducting targeting operations. Under these circumstances, black-box systems are a liability rather than an asset. Moreover, AI methods like advanced neural networks, while promising for other applications, require vast and labeled datasets—often unavailable in tactical contexts. Naïve Bayes methods were also rejected as unsuitable. They exhibited a tendency to assume independent values between variables such as speed, altitude, and heading—an untenable simplification in target analysis. Ultimately, while these methods expedited the mechanics of targeting workflows they failed to capture a critical function: AI-enabled targeting doctrine must codify decision points where human intervention is not just preferred—but required.

While DoD experimentation exercises like Project Convergence are eager to showcase the application of AI technologies, these technologies are not without limitations. For example, current large language models (LLMs), such as Meta's LLaMA, can present a unique risk: They operate via statistical prediction without true comprehension of doctrinal terminology or contextual nuance. If we were to prompt a commercial, off-the-shelf LLM on how to destroy a particular target, this type of AI would not inherently grasp the human understanding of the concept. Believe it or not the targeting effect *destroy* is a complex notion, one filled with contextual relations. Achieving the effect *destroy* beyond simply physical damage also has a time component. *Destroy* also means ensuring the target cannot fulfill its primary function for the remainder of a mission. Such comprehensive understanding requires structured LLM training on doctrinal lexicons, rules-based decision trees, and munitions modeling—all things that are not in these generalist models.

At a fundamental level, incorporating AI into D3A is about optimizing the targeting workflow. In multidomain operations, accelerating sensor-to-shooter kill chains, reducing cognitive burden, and improving commanders' decision-making in contested environments is the goal. By embedding AI where it adds the most value, and ensuring humans remain central at key decision points, the Army can modernize its targeting process while honoring its moral and legal responsibilities. In doing so, it can ensure that D3A remains both fast and just, anchored in human judgment, yet elevated by intelligent machines.

Trust But Verify: US Troops, Artificial Intelligence, and an Uneasy Partnership

Paul Lushenko

Advancements in artificial intelligence have exacerbated the debate surrounding the development of lethal autonomous weapons systems. These killer robots can identify, track, and prosecute targets on their own, without human oversight, and have been used during conflicts in Gaza, Libya, Nagorno-Karabakh, and Ukraine. Critics caution against heightened autonomy in war, citing the potential for abuse that can lead to unintended consequences, including crisis escalation and civilian casualties. Advocates claim the opposite, emphasizing robotic agency in future conflict. They contend that AI-enhanced weapons will encourage human-machine teaming that helps countries maintain lethal overmatch of adversaries while doing so more justly than conventional weapons controlled by humans, particularly because AI is thought to minimize the potential for collateral damage.

Despite rapid advancements in AI and new technologies' growing proliferation on the battlefield, it is unclear what shapes US service members' trust in these technologies. Indeed, scholars have yet to systematically explore this topic, which is puzzling for two reasons.
First, policymakers, academics, and military leaders agree that service members' trust is integral for human-machine teaming, wherein they partner with AI-enhanced military technologies to optimize battlefield performance. Service members are responsible for testing, fielding, and managing the use of AI-enhanced military technologies and, notwithstanding the chain of command, there is no guarantee that they will trust AI. Second, research shows that military capabilities can underdeliver given a lack of user trust, creating vulnerabilities they are designed to overcome, such as prolonged approval for targeting that can have implications for war outcomes.

To determine what shapes service members' trust in human-machine teaming, where they believe that a new AI-enhanced capability will perform as expected, I surveyed senior US miliary officers attending the US Army and Naval War Colleges in Carlisle, Pennsylvania and Newport, Rhode Island, respectively. These are specially selected officers from which the US military will draw its generals and admirals, meaning they will be responsible for adopting emerging capabilities during future conflict. Their attitudes toward human-machine teaming, then, matter.

Contrary to assumptions of service members' automatic trust for AI, which are reflected in emerging warfighting concepts across the US military, I find that service members can be

skeptical of operating alongside AI-enhanced military technologies on the battlefield. Their willingness to partner with these emerging capabilities is shaped by a tightly calibrated set of technical, operational, and oversight considerations. These results provide the first experimental evidence for military attitudes of trust in human-machine teaming and have implications for research, military modernization, and policy.

Three Ingredients of Trust

In 2017, Jacquelyn Schneider and Julia Macdonald published an article in *Foreign Affairs* entitled, "Why Don't Troops Trust Drones: The 'Warm Fuzzy' Problem." They found that soldiers' preferences for drones covaries with operational risk, wherein they trust the use of manned over unmanned aerial vehicles when ground forces are under fire. Though these findings were contested, especially by drone operators themselves, the study benchmarked attitudes of trust among a cross section of the US military. This consisted of those responsible for integrating drones with ground forces, including joint terminal attack controllers and joint fires observers.

We lack comparable evidence relating to service members' trust in partnering with AI on the battlefield. At most, scholars have studied public attitudes of trust toward AI, such as that used in driverless vehicles, police surveillance, or social media, and recommended further study of preferences among the military. While some scholars have researched Australian and Korean soldiers' attitudes toward AI, these studies do not specifically address the trust that is integral to human-machine teaming and are more descriptive than experimental. Thus, scholars cannot draw conclusions for the causal relationships between factors of AI performance and soldiers' trust in partnering with machines during war.

My approach is different. Drawing on extensive research for public attitudes toward drones and AI, I tested how varying nine attributes of AI-enhanced military technologies may shape service members' willingness to partner with these capabilities on the battlefield. Specifically, I used a survey experiment to vary how these capabilities are used, for what outcomes, and with what oversight. To understand service members' trust in partnering with AI in light of these considerations, I administered my survey in October 2023 at the war colleges in Carlisle and Newport. Respondents first read a randomized scenario asking them to consider a hypothetical war in which US military forces use a new AI-enhanced military technology with different attributes. I then asked them to gauge their trust in partnering with the technology using a 5-point Likert scale, where 1 corresponds to "strongly distrust" and 5 corresponds to "strongly trust." I analyzed the data using statistical methods, including by calculating respondents' marginal mean willingness to trust AI-enhanced military technologies given variation in different attributes (Figure 1).

Of note, my sample is not representative of the US military, nor the US Army and Navy. It is a convenience sample that is helpful to draw extremely rare insights into how service members may trust manned-unmanned teaming. My survey returned nearly one hundred high-quality responses, resulting in over eight hundred unique observations, given my within-subject survey design, and offering strong statistical power. This means that respondents received nine randomized scenarios for AI-enhanced military technologies drawn from a pool of thousands of possible attribute pairings, and then gauged their trust for partnering with machines on the battlefield. This pool of officers also offers a hard test for my expectations. They are older, which suggests they may inherently distrust AI-enhanced capabilities. Thus, while I cannot draw sweeping generalizations from my results, I offer the first-ever account for how US service members trust AI-enhanced military technologies given variation in their use, outcomes, and oversight.

Overall, I find that service members question the merits of human-machine teaming in the context I study, which consists of the hypothetical but realistic possibility of a war between the United States and a near-peer adversary. Service members' trust in partnering with AI-enhanced military technologies is based on a combination of three overarching factors. These consist of the technical specifications of machines, their perceived effectiveness, and regulatory oversight.

First, I find that service members' trust for partnering with machines can be heightened when they are used nonlethally, such as for intelligence collection; are not fully autonomous, meaning humans supervise their use; and are highly precise, implying a low false positive rate or minimal target misidentification. Consistent with other respondents, one service member noted that "for strike systems without a human to hold accountable, I want to see precision above 95% accuracy." Indeed, the probability of service members' trust in partnering with machines on the battlefield reduces sharply when they are used autonomously for lethal operations, such as strikes. Research shows that this consideration also reduces public support for killer robots more than other factors.

Second, I find that service members trust partnering with machines when they minimize civilian casualties, maximize the protection of friendly forces, and contribute the most to mission success. This result suggests that service members integrate different moral logics when determining their trust for human-machine teaming. One respondent explained that trust was a function of the "ratio of civilian casualties to US forces saved. I distrust fielding any system that can cause 1K-2K civilian casualties to save 0-2 soldiers. However, I am willing to accept civilian casualties if it saves at least as many US soldiers." Another respondent added that trust was based on the "estimated civilian deaths compared to the estimated military deaths."

These findings are important because they contradict scholars, including Neil Renic and Elke Schwarz at the University of Copenhagen and Queen Mary University, respectively, who argue that emerging technologies have resulted in riskless or post-heroic war. Instead, service members recognize a greater liability to be harmed during war, even when military action is shaped by AI that is thought to enhance troops' protection. My findings also contradict a belief among some experts, such as TX Hammes at the US National Defense University and David Deptula at the Mitchell Institute of Aerospace Studies, that such normative concerns are relatively

inconsequential for the US military's adoption of killer robots. My analysis shows that the probability of service members' trust in human-machine teaming is shaped greatly by moral considerations.

Finally, I find that service members' also trust partnering with AI-enhanced military technologies when their use aligns with international law. Similar to others, one respondent explained that "though I trust US regulation/oversight, it's important for international oversight to ensure compliance to international laws rather than domestic law." This result is consistent regardless of variation in the prospects for mission success and international competition. Even when AI-enhanced military technologies contribute more to the mission, service members demonstrate greater trust in partnering with them when their use comports with international law. A lack of oversight also reduces the probability of service members' trust in human-machine teaming, even when they are informed that perceived adversaries, such as China or Russia, are also adopting AI-enhanced military technologies.

Trust in AI: Hard to Gain, Easy to Lose

These findings suggest that service members' trust is complex and multidimensional. Trust can also be complicated by generational differences across the military ranks. In November 2023, I extended my survey experiment to cadets enrolled in the Reserve Officers' Training Corps across the United States, who are training to become officers. Nearly five hundred cadets completed the survey, resulting in over four thousand unique observations and offering extremely rare insights into how these so-called digital natives perceive human-machine teaming.

I find that cadets trust partnering with AI-enhanced military technologies more so than their future commanders, but not for the reasons scholars typically assume, including misplaced optimism. Crucially, cadets are more willing to partner with AI-enhanced capabilities that are less accurate, implying a greater degree of false positives or target misidentification, particularly when these capabilities contribute the most to mission success.

Together, the results from these related surveys have key research, military modernization, and policy implications. In terms of research, my findings suggest the need for more study, especially in terms of the type of conflict; echelon of use; munition, including nuclear weapons; and target, ranging from personalities to physical structures to networks. As reflected by Erik Lin-Greenberg's research on the escalatory potential of drones, it is possible that these considerations may also shape service members' trust in partnering with AI-enhanced capabilities on the battlefield.

In terms of military modernization, my findings suggest that military leaders should further clarify the warfighting concepts encouraging the development of AI-enhanced military technologies; the doctrine guiding their integration, and across what domain, at what echelon, in what formations, and for what purpose(s); and the policies governing their use. Aligning concepts, doctrine, and policies that govern AI with service members' expectations promises to encourage more trust in human-machine teaming. Finally, if service members' trust in human-machine teaming is, in part, shaped by international oversight, policymakers should explain how US policies on autonomous weapons systems coincide or diverge from international law as well as the norms conditioning their use.

While military leaders claim that human-machine teaming is necessary for success during future wars and often assume service members will trust partnering with AI-enhanced military technologies, my analysis shows that trust is not a foregone conclusion. Trust, like personal integrity, is hard to gain and easy to lose. Amid dizzying developments of AI across sectors and use cases, including within the military, this first evidence provides a convenient roadmap for US political and miliary leaders to enhance service members' trust in human-machine teaming. The costs of not heeding these insights for what shapes service members' trust in human-machine teaming could be the difference between winning or losing in future wars, should analysts' assessments of the implications of AI on the battlefield prove accurate.

Tolstoy's Complaint: Mission Command in the Age of Artificial Intelligence

Theo Lipsky

What will become of battlefield command in the years ahead? This question is at the heart of the US Army's once-in-a-generation reforms now underway. In search of answers, the Army looks to Ukraine. Events there suggest at least two truths. One is that decentralized command, which the US Army calls mission command and claims as its mode, will endure as a virtue. A second is that future commanders will use artificial intelligence to inform every decision—where to go, whom to kill, and whom to save. The recently announced Army Transformation Initiative indicates the Army intends to act on both.

But from these lessons there arises a different dilemma: How can an army at once preserve a culture of decentralized command and integrate artificial intelligence into its every task? Put another way, if at all echelons commanders rely on artificial intelligence to inform decisions, do they not risk just another form of centralization, not at the top, but within an imperfect model? To understand this dilemma and to eventually resolve it, the US Army would do well to look once again to the Ukrainian corner of the map, though this time as a glance backward two centuries, so that it might learn from a young redleg in Crimea named Leo Tolstoy.

What Tolstoy Saw

Before he became a literary titan, Leo Tolstoy was a twenty-something artillery officer. In 1854 he found himself in besieged port of Sevastopol, then under relentless French and British shelling, party to the climax of the Crimean War. When not tending to his battery on the city's perilous Fourth Bastion, Tolstoy wrote dispatches about life under fire for his preferred journal in Saint Petersburg, *The Contemporary*. These dispatches, read across literate Russia for their candor and craft, made Tolstoy famous. They have since been compiled as *The Sebastopol Sketches* and are considered by many to be the first modern war reportage. Their success confirmed for Tolstoy that to write was his life's calling, and when the Crimean War ended, he left military service so that he might do so full-time.

But once a civilian Tolstoy did not leave war behind, at least not as a subject matter. Until he died, he mined his time in uniform for the material of his fiction. In that fiction, most prominently the legendary accounts of the battles of Austerlitz and Borodino found in *War and Peace*, one can easily detect what he thought of command. Tolstoy's contention is that the very idea of command itself is practically a fiction, so tenuous is the relationship between what commanders visualize, describe, and direct and what in fact happens on the battlefield. The worst officers in Tolstoy's stories do great harm by vainly supposing they understand battles at hand when they in fact haven't the faintest idea of what's going on. The best officers are at peace with their inevitable ignorance and rather than fighting it, gamely project a calm that inspires their men. Either way, most officers wander the battlefield, blinded by gun smoke or folds in the earth, only later making up stories to explain what happened, stories others wrongly take as credible witness testimony.

Command or Hallucination?

Students of war may wonder whether Tolstoy was saying anything Carl von Clausewitz had not already said in *On War*, published in 1832. After all, there Clausewitz made famous allowances for the way the unexpected and the small both shape battlefield outcomes, describing their effects as "friction," a term that still enjoys wide use in the US military today. But the friction metaphor itself already hints at one major difference between Clausewitz's understanding of battle and Tolstoy's. For Clausewitz all the things that go sideways in war amount to friction impeding the smooth operation of a machine at work on the battlefield, a machine begotten of an intelligent design and consisting of interlocking parts that fail by exception. As Tolstoy sees it, there is no such machine, except in the imagination of largely ineffectual senior leaders, who, try as they might, cannot realize their designs on the battlefield.

Tolstoy thus differed from Clausewitz by arguing that commanders not only fail to anticipate friction, but outright hallucinate. They see patterns on the battlefield where there are none and causes where there is only coincidence. In *War and Peace*, Pyotr Bagration seeks permission to start the battle at Austerlitz when it is already lost, Moscow burns in 1814 not because Kutuzov ordered it but because the firefighters fled the city, and Russians' masterful knockout flank at Tarutino occurs not in accordance with a preconceived plan but by an accident of logistics. Yet historians and contemporaries alike credit Bagration and Kutuzov for the genius of these events—to say nothing of Napoleon, whom Tolstoy casts as a deluded egoist, "a child, who, holding a couple of strings inside a carriage, thinks he is driving it."

Why then, per Tolstoy, do the commanders and historians credit such plans with unrelated effects? Tolstoy answers this in a typical philosophical passage of *War and Peace*: "The human mind cannot grasp the causes of events in their completeness," but "the desire to find those causes is implanted in the human soul." People, desirous of coherence but unable to espy the many small causes of events, instead see grand things and great men that are not there. Here Tolstoy makes a crucial point—it is not that there are no causes of events, just that the causes are too numerous and obscure for humans to know. These causes Tolstoy called "infinitesimals," and to find them one must "leave aside kings, ministers, and generals" and instead study "the small elements by which the masses are moved."

This is Tolstoy's complaint. He lodged it against the great man theorists of history, then influential, who supposed great men propelled human events through genius and will. But it also can be read as a strong case for mission command, for Tolstoy's account of war suggests that not only is a decentralized command the best sort of command—it is the only authentic command at all. Everything else is illusory. High-echelon commanders' distance from the fight, from the level of the grunt or the kitchen attendant, allows their hallucinations to persist unspoiled by reality far longer than those below them in rank. The leader low to the ground is best positioned to integrate the infinitesimals into an understanding of the battlefield. That integration, as Isaiah Berlin writes in his great Tolstoy essay "The Hedgehog and Fox," is more so "artistic-psychological" work than anything else. And what else are the "mutual trust" and "shared understanding," which Army doctrine deems essential to mission command, but the products of an artful, psychological process?

From Great Man Theory to Great Model Theory

Perhaps no one needs Tolstoy to appreciate mission command. Today American observers see everywhere on the battlefields of Ukraine proof of its wisdom. They credit the Ukrainian armed forces with countering their Russian opponents' numerical and material superiority by employing more dynamic, decentralized command and control, which they liken to the US Army's own style. Others credit Ukrainians' use of artificial intelligence for myriad battlefield functions, and here the Ukrainians are far ahead of the US Army. Calls abound to catch up by integrating artificial intelligence into data-centric command-and-control tools, staff work, and doctrine. The relationship between these two imperatives, to integrate artificial intelligence and preserve mission command, has received less attention.

At first blush, artificial intelligence seems a convincing answer to Tolstoy's complaint. In "The Hedgehog and the Fox" Isaiah Berlin summarized that complaint this way:

Our ignorance is due not to some inherent inaccessibility of the first causes, only their multiplicity, the smallness of the ultimate units, and our own inability to see and hear and remember and record and co-ordinate enough of the available material. Omniscience is in principle possible even to empirical beings, but, of course, in practice unattainable.

Can one come up with a better pitch for artificial intelligence than that? Is not artificial intelligence's alleged value proposition for the commander its ability to integrate all the Tolstoyan infinitesimals, those "ultimate units," then project it, perhaps on a wearable device, for quick reference by the dynamic officer pressed for time by an advancing enemy? Put another way, can't a *great model* deliver on the battlefield what a *great man* couldn't?

The trouble is threefold. Whatever model or computer vision or multimodal system we call "artificial intelligence" and incorporate into a given layer of a command-and-control platform represents something like one mind, but not many minds, so each instance wherein a leader outsources analysis to that artificial intelligence is another instance of centralization. Second, the models we have are disposed to patterns and to hubris, so are more a replication than a departure from the hallucinating commanders Tolstoy so derided. Finally, leaders may reject the evidence of their eyes and ears in deference to artificial intelligence because it enjoys the credibility of dispassionate computation, thereby forgoing precisely the ground-level inputs that Tolstoy pointed out were most important for understanding battle.

Consider the centralization problem. Different models may be in development for different uses across the military, but the widespread fielding of any artificial intelligence–enabled command-and-control system risks proliferating the same model across the operational army. If the purpose of mission command were strictly to hasten battlefield decisions by replicating the mind of a higher command within junior leaders, then the threat of centralization would be irrelevant because artificial intelligence would render mission command obsolete. But Army Doctrinal Pamphlet 6-0 lists as mission command's purpose also the levering of "subordinate ingenuity"—something that centralization denies. In aggregate one risks giving every user the exact same coach, if not the exact same commander, however brilliant that coach or commander might be.

Such a universal coach, like a universal compass or rifle, might not be so bad, were it not for the tendency of that universal coach to hallucinate. That large language models make things up and then confidently present them as truth is not news, but it is also not going away. Nor is those models' basic function, which is to seek patterns and then extend them. Computer vision likewise produces false positives. This "illusion of thinking," to paraphrase recent research, severely limits the capacity of artificial intelligence to tackle novel problems or process novel environments. Tolstoy observes that during the invasion of Russia "a war began which did not follow any previous traditions of war," yet Napoleon "did not cease to complain . . . that the war was being carried on contrary to all the rules—as if there were any rules for killing people." In this way Tolstoy ascribes Napoleon's disastrous following Borodino precisely to the sort of error artificial intelligence is prone to make—the faulty assumption that the rules that once applied extend forward mechanically. There is thus little difference between the sort of prediction for which models are trained and the picture of Napoleon in *War and Peace* on the eve of his arrival in Moscow. He imagined a victory that the data on which he had trained indicated he ought expect but that ultimately eludes him.

Such hallucinations are compounded by models' systemic overconfidence. Research suggests that, like immature officers, models prefer to confidently proffer an answer than confess they just do not know. It is then not hard to imagine artificial intelligence processing incomplete reports of enemy behavior on the battlefield, deciding that the behavior conforms to a pattern, filling in gaps the observed data leaves, then confidently predicting an enemy course of action disproven by what a sergeant on the ground is seeing. It is similarly not hard to imagine a commander directing, at the suggestion of an artificial intelligence model, the creation of an engagement area anchored to hallucinated terrain or queued by a nonexistent enemy patrol. In the aggregate, artificial intelligence might effectively imagine entire scenarios like the ones on which it was trained playing out on a battlefield where it can detect little more than the distant, detonating pop of an explosive-laden drone.

To be fair, uniformed advocates of artificial intelligence have said explicitly that no one wants to replace human judgment. Often those advocates speak instead of artificial intelligence informing, enhancing, enabling, or otherwise making more efficient human commanders. Besides, any young soldier will point out that human commanders make all the same mistakes. Officers need no help from machines to spook at a nonexistent enemy or to design boneheaded engagement areas. So what's the big deal with using artificial intelligence?

The issue is precisely that we regard artificial intelligence as more than human and so show it a deference researchers call "automation bias." It's all but laughable today to ascribe to any human the genius for seeing through war's complexity that great man theorists once vested in Napoleon. But now many invest similar faith in the genius of artificial intelligence. Sam Altman of OpenAI refers to his project as the creation of "superintelligence." How much daylight is there between the concept of superintelligence and the concept of the great man? We thus risk treating artificial intelligence as the Napoleon that Napoleon could not be, the genius integrator of infinitesimals, the protagonist of the histories that Tolstoy so effectively demolished in *War and Peace*. And if we regard artificial intelligence as the great man of the history, can we expect a young lieutenant to resist its recommendations?

What Is to Be Done?

Artificial intelligence, in its many forms, is here to stay. The Army cannot afford in this interwar moment a Luddite reflex. It must integrate artificial intelligence into its operations. Anybody who has attempted to forecast when a brigade will be ready for war or when a battalion will need fuel resupply or when a soldier will need a dental checkup knows how much there is to be gained from narrow artificial intelligence, which promises to gain immense efficiencies in high-iteration,

structured, context-independent tasks. Initiatives like Next Generation Command and Control promise as much. But the risks to mission command posed by artificial intelligence are sizable. Tolstoy's complaint is of great use to the Army as it seeks to understand and mitigate those risks.

The first way to mitigate the risk artificial intelligence poses to mission command is to limit the use of it those high-volume, simple tasks. Artificial intelligence is ill-suited for low-volume, highly complex, context-dependent, deeply human endeavors—a good description of warfare—and so its role in campaign design, tactical planning, the analysis of the enemy, and the leadership of soldiers should be small. Its use in such endeavors is limited to expediting calculations of the small inputs human judgment requires. This notion of human-machine teaming in war is not new (it has been explored well by others, including Major Amanda Collazzo via the Modern War Institute). But amid excitement for it, the Army risks forgetting that it must carefully draw and jealously guard the boundary between human and machine. It must do so not only for ethical reasons, but because, as Tolstoy showed to such effect, command in battle humbles the algorithmic mind—of man or machine. Put in Berlin's terms, command remains "artistic-psychological" work, and that work, even now, remains human work. Such caution does not require a ban on machine learning and artificial intelligence in simulations or wargames, which would be self-sabotage, but it does require that officers check any temptation to outsource the authorship of campaigns or orders to a model—something which sounds obvious now, but soon may not.

The second way is to program into the instruction of Army leaders a healthy skepticism of artificial intelligence. This might be done first by splitting the instruction of students into analog and artificial intelligence–enabled segments, not unlike training mortarmen to plan fire missions with a plotting board as well as a ballistic computer. Officers must first learn to write plans and direct their execution without aid before incorporating artificial intelligence into the process. Their ability to do so must be regularly recertified throughout their careers. Classes on machine learning that highlight the dependency of models on data quality must complement classes on intelligence preparation of the battlefield. Curriculum designers will rightly point out that curricula are already overstuffed, but if artificial intelligence–enabled command and control is as revolutionary as its proponents suggest, it demands a commensurate change in the way we instruct our commanders.

The third way to mitigate the risks posed is to program the same skepticism of artificial intelligence into training. When George Marshall led the Infantry School during the interwar years, he and fellow instructor Joseph Stilwell forced students out of the classroom and into the field for unscripted exercises, providing them bad maps so as to simulate the unpredictability of combat. Following their example, the Army should deliberately equip leaders during field exercises and wargames with hallucinatory models. Those leaders should be evaluated on their ability to recognize when the battlefield imagined by their artificial intelligence–enabled command-and-control platforms and the battlefield they see before them differ. And when training checklists require that for a unit to be fully certified in a task it must perform that task under dynamic, degraded conditions, "degraded" must come to include hallucinatory or inoperable artificial intelligence.

Even then, Army leaders must never forget what Tolstoy teaches us: that command is a contingent, human endeavor. Often battles represent idiosyncratic problems of their own, liable to defy patterns. Well-trained young leaders' proximity to those problems is an asset rather than a liability. For that proximity they can spot on the battlefield infinitesimally small things that great data ingests cannot capture. A philosophy of mission command, however fickle and at times frustrating, best accommodates the insights that arise from that proximity. Only then can the Army see war's Tolstoyan infinitesimals through the gun smoke and have any hope of integrating them

The End of Audacity? Artificial Intelligence and the Future of Command

Antonio Salinas and David V. Gioe

Some of military history's most decisive victories didn't come from perfect planning but from bold risks whose chances of success were so slim that no modern algorithm would ever recommend them. Still, we remember commanders who seem to have achieved the impossible at considerable indifference to overwhelming odds. Call it the *hold my beer* moment: Miltiades at Marathon, Chamberlain's bayonet charge at Little Round Top, or Eisenhower launching D-Day through a narrow weather window. These were calculated risks that challenged the odds, surprised the enemy, and changed the course of history.

Today, as people increasingly rely on artificial intelligence and decision-support systems to guide their choices, they form courses of action based on extensive data. However, machines and large language models are designed to favor statistical methods with higher success rates over Clausewitzian calculations of chance, moral forces, and human instinct that seek to seize fleeting opportunities. The risk for military commanders is that, in the name of harnessing AI, we might lose the willingness to make bold, high-risk decisions in the moment, especially if AI recommends otherwise. If we entrust war to the machine's logic, we may win battles of efficiency but lose the wars of will. For all its remarkable capabilities, artificial intelligence lacks the human will to dare.

There is now potential for the military to increasingly use AI tools like large language models (LLMs) to quickly and

effectively integrate intelligence and to model courses of action for commanders. An LLM-embedded military decision-making process can identify, analyze, and integrate vast amounts of data that far surpasses the capabilities of our modern planners. Indeed, thinking and writing without AI tools is like shooting a rifle without a scope. While the modern AI-enabled staff can develop plans and courses of action with speed and understanding far beyond what our experienced officers staring at a map can do, we wonder if there is an accompanying risk to audacity. But the real question is not whether AI will replace audacity. It is whether militaries will design, integrate, and culturally absorb these tools in ways that preserve (or could undermine) the human capacity for bold judgment under uncertainty.

Who Dares Wins

AI can never calculate the primal horror, fear, and chaos of battle. Combat consistently demands decisions under extreme uncertainty and cognitive stress. Actions such as participating in a room-clearing stack, advancing across exposed terrain, or closing with an adversary require decisive action when calculation alone is insufficient. Carl von Clausewitz reminds us that boldness is the very steel that gives the sword its edge and brilliance.

Audacity is not the same as recklessness. Risk calculus has always mattered in military decision-making. Clausewitz famously emphasized chance, friction, and moral forces—factors that are hard to calculate but cannot be ignored. Many of history's audacious battlefield decisions were anything but impulsive. The Normandy invasion was preceded by years of planning, intelligence collection, and long-term deception operations. Eisenhower's fateful decision to go on June 6, 1944 was bold precisely because it rested on an informed appreciation of uncertainty, not ignorance of it. The gamble lay not in disregarding analysis, but in accepting its limits.

Marathon: Betting Everything on Shock

While there is no sure way to win a battle when outnumbered, there is one guaranteed way to lose: Do nothing. In 490 BCE, a small force of ten thousand Athenians and their allies faced a Persian army of over twenty thousand at the Battle of Marathon. Any gambler would have bet on the Persians in this fight, and perhaps AI would have too. They had numerical superiority, advantageous terrain, and momentum at their backs, having just won a victory over another Greek city-state called Eretria, where they enslaved its population. The Persians expected yet another easy win. At first, the Athenians stayed in their camp on high ground, watching the massive Persian force assemble on the beach in front of their position.

Yet, as the battle commenced, the Athenians chose not to hold the hill and fight on the defensive. Instead, they did the unthinkable: They left the high ground and charged a numerically superior force.

The charge of the Athenians seized victory from the jaws of defeat and has resonated throughout history, remaining alive in university and military academy classrooms. This charge was perhaps one of the most famous gambits that defied conventional wisdom in Western military history. The Greeks quickly closed the distance with their Persian enemies on the plain of Marathon and collided with their lightly armored ranks. The Persians were not given the chance to wage the battle they preferred. They were not allowed to rain arrows down on the Greeks or use cavalry. Instead, the Greeks fought this battle at close quarters. The Athenians unleashed a human tidal wave of bronze and sinew. The charge of the Athenians amplified the existing strength of the phalanx, the advantage of hoplite equipment, and the initiative of the Greek general Miltiades.

On the display of any AI decision tool, the Athenians' chosen course of action would have been painted red—high risk, low probability, avoid. But on the field of battle, it worked. The sudden shock broke the Persian line, and the Persians fled to their ships. The gamble saved Athens, preserved Greek independence, and indirectly set the stage for the rise of Western democracy.

Little Round Top: Bayonets Against the Odds

Fast forward to July 2, 1863, the second day of the Battle of Gettysburg. On the Union Army's far left flank, Colonel Joshua Lawrence Chamberlain and the 20th Maine were ordered to hold Little Round Top "at all hazards." By this point in the Civil War, the 20th Maine was not a full regiment. After months of hard campaigning, sickness, and casualties, Chamberlain's ranks had been badly depleted—barely 350 men were prepared to defend the rocky hill that anchored the entire Union line. Facing them were waves of Confederate assaults from seasoned troops who were veterans hardened by years of combat.

As the day dragged on, the 20th Maine fought through heat, exhaustion, smoke, and chaos. The Union soldiers repelled charge after charge, firing until their ammunition nearly ran out. It seemed that the situation was hopeless. Reason—and any modern algorithmic decision-support system—would have recommended withdrawal. Their line was thin, their flank exposed, and their cartridges almost depleted. But Chamberlain understood what machines can't: the intangible factors of momentum and morale. A retreat here could unravel the entire Union position. He grasped in that moment that the only way to hold was to attack. When the next Confederate charge climbed the slope, Chamberlain gave an audacious order: Fix bayonets.

With a shout the remnants of the 20th Maine surged forward in a wheeling charge, crashing into the stunned Confederate lines. So unexpected was this counterattack that the commander of the 15th Alabama, Lieutenant Colonel William C. Oates, believed Chamberlain must have received reinforcements. In truth, no reinforcements had come; Chamberlin and the 20th Maine were out of everything except raw human courage. That daring charge broke the Confederate assault, took dozens of prisoners, and protected the Union flank. It was a bold gamble that no algorithm would ever support. Yet that single human decision, made in chaos and courage, helped tip the scales not just of a battle, but of a war that changed America for the better.

D-Day: Through the Weather Window

In the days leading up to the D-Day invasion, Allied commanders studied meteorological charts filled with bad

news. The weather over the English Channel was stormy and unpredictable—high winds, low clouds, and heavy seas battered the invasion fleet's staging areas. Landing 156,000 troops, thousands of vehicles, and mountains of equipment under such conditions seemed impossible. Logic suggested a delay.

The safer call, and indeed one that many staff officers urged, was to wait for a better window. A decision-support system that assessed probabilities probably would have recommended the same. But General Dwight D. Eisenhower understood something no machine could quantify: the intangible costs of hesitation on his armada. Secrecy and the deception campaigns were already stretched to their limits. Each day of delay gave the Germans more time to reinforce beaches, mine approaches, and strengthen defenses. Waiting for perfect conditions could mean missing the only fleeting chance for surprise.

The Germans, for their part, were confident that no invasion was imminent. Their meteorologists, cut off from Atlantic weather data, predicted that the storm would last for days. Field Marshal Erwin Rommel left his headquarters to celebrate his wife's birthday, convinced that the Channel seas made invasion impossible. So confident were the defenders that German panzer divisions were placed under strict higher command and could not move without the approval of Hitler, who, asleep at his headquarters, was not woken until hours after the landings began. But the Allies, aided by Atlantic weather data the Germans lacked, detected a narrow thirty-six-hour break in the tempest. Eisenhower seized it. At 4:15 a.m. on June 5, after long silence and visible strain, he said, "Okay, let's go."

The gamble paid off. The storm still raged, but the Germans were unprepared; their defenses were manned at half strength, their armor still held in reserve. A commander overly reliant on AI's calculations would have waited for clear skies. Eisenhower read the chaos and chose audacity over caution. That single leap through the storm changed the fate of the world.

Clausewitz, Chance, and Moral Forces

Clausewitz stated that war is influenced by violence, chance, and reason, a remarkable trinity that connects the passions of the people, the unpredictability of the commander, and the political goals of the state. In contrast, AI, by design, will try to tame chance and make reason dominant. It will smooth out the volatility of human emotion, compress uncertainty with better data, and offer courses of action that minimize risk. In doing so, it could fundamentally alter the balance of the trinity. The trouble is that the irrational element—the willingness to accept great risk to win—has often been the spark that turns a stalemate into a victory. Machines will calculate and weigh probabilities, but they cannot recognize the fleeting moment when risk becomes opportunity.

This distinction matters because contemporary debates about AI often collapse judgment and calculation into a false binary. Decision-support systems do not make decisions. They structure information, generate options, and illuminate trade-offs. Whether they encourage caution or enable boldness depends on how commanders use them, and how institutions reward or punish risk.

In some cases, AI may actually *enable* audacity. Better situational awareness, faster data fusion, and improved logistics forecasting can give commanders the confidence to accept risks they might otherwise avoid. A clearer understanding of adversary vulnerabilities or operational constraints can expand, rather than narrow, the menu of feasible options. Historically, uncertainty has not always bred boldness; it has often produced paralysis. Clausewitz might have recognized AI as a tool, but he would warn against letting it reshape war into a purely rational exercise devoid of passion.

The Risk of Algorithmic Caution

Modern militaries introduced analytical tools, staff processes, and decision aids in part because human judgment is fallible, prone to overconfidence, groupthink, and wishful thinking. AI is the latest iteration of a long effort to discipline those weaknesses. The risk is not that militaries will become too rational, but there is a legitimate danger of automation bias. Humans tend to defer to systems that appear authoritative, especially under time pressure. If decision-support tools consistently privilege probabilistic success, minimized losses, or institutional risk aversion, commanders may find it psychologically and professionally harder to override them, even when circumstances demand it. Over time, this can reshape organizational norms, subtly redefining what *reasonable* risk looks like.

Decision-support systems will be highly effective for specific tasks: quickly analyzing battlefield data, optimizing logistics and force deployment, and simulating likely enemy reactions. However, they will also overlook certain critical factors by potentially underestimating the psychological effects of bold actions, overvaluing numerical safety, and failing to grasp aspects that refuse to present as data, like morale, willpower, and fear. In other words, they will be bad at recognizing and exploiting moments when audacity will be rewarded, not because the probabilities are wrong, but because such moments defy calculation.

The problem, then, is not AI per se, but how militaries encode risk into their tools and cultures. Algorithms are not neutral. They reflect the assumptions, priorities, and incentives of the institutions or personnel that design and deploy them. A force that prizes force protection above mission accomplishment will build different systems than one that rewards initiative and accepts calculated losses. Technology will amplify those preferences, not replace them.

This is where our historical analogies require caveats. It is tempting to contrast heroic gambles that succeeded with hypothetical algorithmic caution that would have prevented them. But history is also littered with audacious failures as well as successes: Gallipoli in World War I and Operations Market Garden and Barbarossa in World War II. These great gambles all involved boldness and commander's assumptions exceeding realistic assessment that would not have required AI to ascertain. Survivorship bias and the martial appreciation for heroism and glory can skew perceptions of the results of audacity.

If commanders become accustomed to relying on the machine, their risk tolerance may decrease, especially if an institutional postmortem after a failed operation cites AI's probabilities that counseled caution as a reason to second-guess the commander's boldness. Over time, armed forces might shift toward strategies that are more predictable, and predictable opponents are easier to defeat. The solution isn't to reject AI. Its ability to gather and process information quickly is a gift no modern commander should overlook. However, we must intentionally shape doctrine, training, and command culture so that AI recommendations are considered, rather than replaced, by human judgment.

Preserving Audacity in the Age of Algorithms

The deeper issue is command responsibility. No algorithm bears moral or strategic accountability for failure. That burden rests with human commanders and political leaders. If institutions begin to treat AI recommendations as default answers rather than inputs to judgment, responsibility becomes blurred. Decisions can start to feel *validated by systems* rather than *owned by commanders*. In such environments, audacity does not disappear, but it might become subtly institutionally discouraged. This dynamic is already visible in fields beyond the military profession. Financial markets, medical diagnostics, and aviation all wrestle with the tension between automation and professional judgment. In each case, the most resilient systems are those that deliberately preserve human override, cultivate skepticism toward automated outputs, and train professionals to understand not just what systems recommend, but why.

For militaries, this implies several practical imperatives. First, AI systems should be designed to surface uncertainty, not obscure it. As a good intelligence professional might do, the AI should highlight confidence intervals, assumptions, and data gaps. Such programmed transparency would reinforce the reality that judgment is still required. Second, military education should explicitly address how to *disagree with machines*. Teaching officers when and how to override decision aids is as important as teaching them how to use them. Third, organizational incentives matter. If promotion, evaluation, and after-action processes punish deviation from algorithmic recommendations (even when outcomes justify it), commanders will learn to conform. Conversely, if institutions reward informed risk-taking and honest failure, audacity remains possible.

Technology cannot compensate for cultures that fear command responsibility or algorithmically reduce it. War remains a profoundly human enterprise, shaped by will, perception, and emotion as much as by calculation, plans, and training. Clausewitz's friction has not disappeared; it has migrated into new dimensions, including cyber, information, and machine-human interaction. In the AI age, resisting that pull toward machine-based recommendations will require deliberate effort. AI may assist our ability to reason, but it cannot feel the tremor in the chest before a charge, the weight of duty that defies odds, or the surge of courage. Audacity will not vanish because algorithms exist. It will vanish only if institutions allow judgment to atrophy behind the appearance of optimization. The challenge is not to choose between AI and audacity, but to ensure that one does not quietly crowd out the other.

The history of Marathon, Little Round Top, and D-Day shows us that some victories come only to those willing to take a fateful plunge. AI will change the character of war, but it must not strip away its art. The soul of victory has always belonged to those who dare.

Hold my beer, indeed.

Unmanned Systems

Battlefield Drones and the Accelerating Autonomous Arms Race in Ukraine

Samuel Bendett and David Kirichenko

Since Russia's full-scale invasion of Ukraine in February 2022, the war there has been impacted by attritable, cheap drones and rapidly growing roster of unmanned and robotic systems. Collectively, these technologies are redefining how military forces can wage modern warfare. With both sides in this war rushing to secure a technological advantage, the Ukrainian battlefield is transforming into a clash between conventional forces backed by a growing number of autonomous and remote-controlled systems. Both Ukraine and Russia have steadily poured more and more resources into developing this technology in a bid to stay a step ahead of the adversary.

Ukraine's battlefield experience reflects a shift toward unmanned systems that augment or attempt to replace human operators in the most dangerous missions, and against an enemy willing to send more and more manpower into large-scale frontal assaults. After so many autonomous and robotic systems were fielded over the past three years by Kyiv's forces, Ukrainian officials started to describe their country as a "war lab for the future"—highlighting for allies and partners that, because these technologies will have a significant impact on warfare going forward, the ongoing combat in Ukraine offers the best environment for continuous testing, evaluation, and refinement of such systems. Many companies across Europe and the United States have tested their drones and other systems in Ukraine. At this point in the conflict, these companies are striving to gain "battle-tested in Ukraine" credentials for their products.

For example, US defense tech company Anduril recently started selling its new autonomous drones after successful tests carried out in Ukraine in October 2024. Ukrainian and Western drone manufacturers have started partnering more closely both on drones and on certain types of AI development. The US military is seeking to speed up the deployment of cheap autonomous systems through its Replicator program, and is also working closely with the private sector to test systems and technologies in Ukraine that can then be potentially used in future conflicts.

Recently, US Army Chief of Staff General Randy George noted that the Ukraine war "has demonstrated the value of small, attritable drones on the battlefield." This combat application of relatively inexpensive platforms has provided the Pentagon with an opportunity to see how integrating cutting-edge software with scalable drone technology can proceed across the US Department of Defense, drawing lessons from the Russia-Ukraine war as it prepares for potential future conflicts, including with China.

In December 2024, for the first time, Ukrainian forces successfully carried out an attack on Russian positions using only ground and first-person view drones, further evolving how Ukraine is leveraging unmanned technology on the battlefield. According to Sergeant Volodymyr Dehtiarov of the Khartiia Brigade involved in this attack, dozens of robotic and unmanned systems, including machine-gun-equipped ground drones and kamikaze first-person view aerial drones, were deployed near Lyptsi, north of Kharkiv. While these were remote-controlled systems that still required a large human complement to operate them, this is the first step in the process of Ukraine gradually working to deploy more combat robots and eventually bring more autonomous systems to the battlefield. Ukraine also previously used a ground robot in an assault on a Russian trench in Kursk Oblast, in September 2024, with numerous other examples of such systems being rapidly built and fielded for combat. In many ways, Ukraine has no choice but to maximize its use of technology, as the manpower disparity between Ukraine and Russia is still significant along the eight-hundred-mile front line of the war.

While technological developments have proceeded at a very rapid pace in this war, it also became clear that systematizing the combined research, development, testing, evaluation, and use of different systems by different units across the entire force was crucial. Therefore, in February 2024, Ukraine's president Volodymyr Zelenskyy signed a decree to establish the national Unmanned Systems Forces, with Colonel Vadym Sukharevskyi appointed as commander in June 2024. In December 2024, the Russian military followed up by announcing that it was establishing an unmanned systems branch to better integrate its forces' use of autonomous and robotic technologies, and to make sure that lessons and tactics from combat in Ukraine can be absorbed and codified by different military branches.

Both countries also claim multiple AI developments for their respective militaries, in drones as well as in other battlefield systems and tactical applications. Three years into its war against Russian aggression, Ukraine has led the way in conceptualizing large-scale development and application of different unmanned systems and AI technologies across domains and different mission sets. In 2025, Ukraine is expected to field AI-enabled drone swarms and massive numbers of ground vehicles to counter Russian forces. As one Ukrainian official put it: "We count people, and we want our people to be as far from the front line as we can."

Ukraine's private sector has stepped up to accelerate the development of autonomous and robotic technologies for enhanced targeting capabilities, with companies like TAF Drones leading the way, aided by the Brave1 organization, a coordination platform established by Ukraine's government playing an important role in helping the private sector. Ukraine's plan is to ensure AI-powered combat drones can ensure the nation's advantage over the Russian force on the battlefield. The Russian military claims the same for its military AI research and application in this war.

For example, Russian Defense Minister Andrei Belousov stated in October 2024 that AI-powered drones are playing a pivotal role on the battlefield in Ukraine, though he did not elaborate further. To better understand how different types of robotic and autonomous systems are used in Ukraine combat, the Russian Ministry of Defense launched the Rubicon Center in August 2024 to help systematize lessons from Ukraine, including the development and application of AI. This initiative is likely to be the epicenter for Russia's formation of its planned unmanned systems branch. Russian president Vladimir Putin also announced that Russia is increasing military drone production to approximately 1.4 million in 2024, aiming to stay abreast of Ukraine's own rapid and large-scale drone manufacturing.

Both Ukrainian and Russian forces prioritize minimizing drone operator involvement to protect trained assets in a complex combat environment. Ukraine's survival-driven focus often outweighs ethical concerns tied to lethal autonomous weapon systems. Meanwhile, despite recent announcements of AI-enabled combat drones already used against Ukraine, Russia's military AI likely mainly supports data analysis and rapid decision-making. For example, In November 2024, the Russia-allied Donetsk People's Republic claimed that its "Donbass Dome" airspace defense and electronic warfare system evaluates different types of information from multitudes of sources to evaluate incoming threats, allegedly done with the help of artificial intelligence algorithms. The evaluated data is transmitted to the military and law enforcement for follow-on actions.

Considering the Russian military's attempt at making sense of the Ukrainian battlefield, such data analysis efforts are likely taking place across different systems, though public information on their overall effectiveness is relatively scarce. Similar efforts exist across the Russian defense sector, with a subsidiary of national industrial giant Rostec claiming in 2024 the development of a neural network for optical drone detectors, which allegedly allows for increasing their detection range by 40 percent.

On the other side of the war, Ukrainian officials are on record noting the need for tens of thousands of uncrewed robotic ground vehicles in 2025 for combat and logistics missions. These officials also noted that Ukrainian forces have been using dozens of domestically made AI-augmented systems to enable aerial drones to reach targets on the battlefield without being piloted and remain effective in areas protected by extensive jamming. At this point in the war, there are around ten Ukrainian companies competing in state procurements to offer AI products.

Ukrainian officials have stated that in 2025, more autonomous drones with AI targeting will arrive on the battlefield, potentially making way for "real drone swarm uses." Ukraine's efforts to use AI on the battlefield are aided by willing partners, such as the Germany-based Helsing AI firm. In December 2024, Helsing announced that the first few hundred of almost four thousand of its AI-equipped HX-2 Karma unmanned aerial vehicles earmarked for Ukraine were set to be delivered to the Ukrainian front. Apparently, HX-2 is immune to electronic warfare countermeasures via its ability to search for, reidentify, and engage targets without a signal or a continuous data connection, while allowing a human operator to stay in or on the loop for critical decisions.

Russian technical experts already acknowledge that "autonomous flying robots"—drones with artificial intelligence that determine their own targets—are already used in combat and apparently "kill" people, though they usually don't provide technical specifications for such claims. It is likely that such developments indicate a more limited AI role in aerial drones, such as the terminal guidance and image recognition that allow drones to fly autonomously to designated targets once the human operator has approved strikes on said targets.

While on the receiving end of Ukraine's increasing AI and autonomy use, many Russian experts express concerns that the pace of AI-enabled military developments could get out of control, thus requiring global regulation "in the interests of all humanity," while also noting the difficulty of banning the development of AI for military purposes while the outcomes of wars hang in the balance and national interests are at stake. Still, Russian military experts, such as those writing in key military publications like *Arsenal Otechestva*, believe in AI's potential in military applications. These experts highlight its ability to enhance system autonomy, improve tactical decision-making, enable real-time operational support in combat zones, reduce crew risks, and decrease uncertainty through rapid processing of large, unstructured data.

With Russia determined to fight until Ukraine is conquered, and Ukraine resolute in defending its freedom, the technological arms race in this war continues to accelerate. Each month in this protracted war brings new technological developments and achievements, with the innovation cycle continuously driven forward by new technologies that are either copied or countered by the adversary, sparking a fresh round of innovation to achieve the next breakthrough.

Ukraine's Western supporters are closely monitoring how such technologies are developed and fielded in combat. Retired Army General Mark Milley, former chairman of the Joint

Chiefs of Staff, has predicted that within the next ten to fifteen years, up to one-third of the US military could consist of robotic systems, an assessment likely informed by observations of technologies fielded in the Ukraine war. To be sure, certain systems in use by both Ukrainian and Russian forces can function more effectively than others on a battlefield teeming with countermeasures, but the sum total of different autonomous, robotic, and unmanned technologies used in the past three years demonstrates the potential for rapid, large-scale fielding. Both Ukraine and Russia are continuously accelerating their development of different types of battlefield drones and robotic systems, driven by the need for precision, mass employment to overwhelm the adversary, resilience against countermeasures, and reducing risks to human lives. These advancements are impacting the battlefield at the tactical and operational levels and are shaping how future warfare may be conducted.

What Iran's Drone Attack Portends for the Future of Warfare

Joshua A. Schwartz

Iran's attack against Israel on April 14 was historic—it marked the first time that Iran has directly struck Israeli territory from its own soil despite decades of tensions and shadow conflict. Iran utilized around 170 drones in the operation, making it one of the largest drone attacks in history—possibly *the* largest. As such, the attack epitomizes the increasing reliance on remote, uninhabited systems in modern warfare.

Aerial drones and other types of uninhabited vehicles are undoubtedly key to the future of conflict, but Iran's attack demonstrates that the current generation of these systems have crucial weaknesses that limit their effectiveness on the battlefield against sophisticated adversaries. In particular, drones are highly susceptible to air defense and thus often do not reach their intended targets. However, Iran's large-scale use of drones against Israel also illustrates how the *military* deficiencies of these systems can be leveraged to achieve two higher-order, strategic *political* goals—limiting escalation and maintaining a strong reputation for resolve.

Defense Is Stronger Than You Might Think

The only thing more striking than the large quantity of drones Iran used in its attack against Israel was the number of those drones that were shot down by Israel and other countries. According to Israeli estimates, over 99 percent of all Iranian weapons used in the attack were intercepted before reaching their targets—including all 170 drones. In part, this reflects the sophistication of Israel's air defense capabilities and the abilities of the many other countries that helped Israel destroy these drones. But it also highlights something broader—the generally high susceptibility of drones to air defense compared to more traditional inhabited aircraft.

There are at least three reasons uninhabited aircraft are typically easier to shoot down than their inhabited counterparts. First, current-generation drones tend to fly much slower. For example, Iran's Shahed-136 drones, which were used in the attack against Israel, can only fly a maximum speed of around 115 miles per hour. By contrast, Iran's inventory of MiG-29 inhabited aircraft, which it acquired decades ago in the early 1990s, have maximum speeds closer to 1,500 miles per hour. The slow speed of uninhabited aircraft has helped enable Ukraine to shoot down Russian drones (many provided by Iran) with even unsophisticated air defense tools like machine guns.

Second, today's drones tend to have only limited countermeasures they can deploy to protect themselves against air defense systems. For instance, they typically do not carry chaff or flares, which can be used to confuse air defense missiles. Compared to inhabited aircraft, military-grade drones (such as the Shahed or the Turkish-built Bayraktar TB-2 drone used by Ukraine) usually have quite limited maneuverability. This weakness, which does not apply to small quadcopters, makes it harder for drones to evade air defense missiles by executing sudden rolls and turns.

Third, the signals that enable communication between a pilot and a drone can be jammed. This is one crucial defense tool Russia and Ukraine have been using to down each other's drones. It is also a tactic Israel deployed to disrupt the Iranian attack.

Of course, the cat-and-mouse game between drones and air defense will spur future innovations that could make uninhabited aerial vehicles less suspectable to being shot down. For example, drones can be designed to fly at faster speeds, carry more sophisticated countermeasures to air defense systems, and operate autonomously if communication links with pilots are severed. Furthermore, even existing systems do have at least one potential advantage over the defense: shooting down cheap drones that cost just tens of thousands of dollars with expensive air defense assets that can cost hundreds of thousands of dollars or more can bleed the financial resources of a country over time. Israel's defense likely cost more than Iran's offensive.

Nevertheless, the high vulnerability of most current-era drones to air defense can help explain why all of the Iranian drones were shot down and failed to reach their intended targets. It also explains why the attrition rates of Ukrainian and Russian drones are similarly high, with Ukraine losing as many as ten thousand drones per month. As one Ukrainian air force pilot said, relatively high-end and expensive Turkish TB-2 drones "were very useful and important in the very first days [of the war] . . . but now that [the Russians have] built up good air defenses, they're almost useless."

While many types of drones—especially cheaper, attritable systems—are indeed extremely useful on the battlefield, arguments that drones provide a significant advantage to the offense over the defense are at least somewhat overstated. Countries should thus not consider drones as a panacea, especially when operating against adversaries with relatively advanced air defense systems.

Turning a Weakness Into a Strength

Iran's attack was not particularly successful from a military or operational perspective in that it failed to hit and inflict significant damage on almost all of its targets. However, it may have been successful from a political perspective in that it helped enable Iran to achieve two of its strategic goals: limiting escalation and maintaining a high reputation for resolve.

Since the devastating Hamas attack against Israel on October 7, it has been clear that Iran has little interest in igniting a wider war in the Middle East. On October 29, Iranian Foreign Minister Hossein Amir-Abdollahian publicly said, "We don't want this war to spread out." In private, Iranian Supreme Leader Ayatollah Ali Khamenei reportedly ordered his military subordinates to adopt a policy of "strategic patience" to avoid escalation. Iran's deeds also match its words (at least to some extent). For example, Iran has reportedly urged its chief proxy, Hezbollah, to exercise restraint and refrain from launching significant attacks against Israeli territory. Attempting to limit escalation is rational given that Israel is more capable militarily than Iran. A wider war would also risk the United States' direct involvement in military operations against the Islamic Republic, which is surely a dynamic the supreme leader wishes to avoid.

The use of drones and other remote weapons, such as missiles, helps Iran achieve its goal of limiting escalation with Israel and the United States. Precisely because the Iranian drones failed to hit their mark and cause significant destruction, Israel and the United States were under less pressure to respond forcefully in ways that might raise the risk of a wider war.

In accordance with the logic, President Joe Biden urged Israel not retaliate and told Israeli Prime Minister Netanyahu, "You got a win. Take the win." Israel chose not to fully take Biden's advice and instead did retaliate against Iran by conducting its own strike against an air defense system near the Iranian city of Isfahan. However, the attack was small, was limited in nature, and appears to have caused little major damage. In fact, Israel had initially planned on a more significant counterattack against Iran, but ultimately settled on a smaller-scale retaliation due to foreign pressure and the ineffectiveness of Iran's attack. Israel also has incentives to avoid major escalation given that a wider conflict would put it in the precarious position of having to fight a three-front war—against Hamas in Gaza, Iran to the east, and Iran's proxy Hezbollah to the north in Lebanon.

Iran's reaction to the Israeli counterattack has been muted, indicating that a de-escalation of the immediate crisis is probable. Iran's use of drones and other remote systems in the initial attack against Israel is one reason why the Iranian regime was under less pressure to respond forcefully to Israeli retaliation, which could have led to an escalation spiral of attacks and counterattacks. As demonstrated in experimental wargames conducted by MIT professor Erik Lin-Greenberg that presented variable scenarios to individuals with military experience, the shooting down of a drone is less likely to lead to escalation because it does not put at risk a human life. Iran learned this lesson firsthand following its destruction of an expensive American reconnaissance drone in 2019. While President Donald Trump nearly authorized a direct retaliatory attack against Iran, he ultimately changed his mind and noted such a strike is "not proportionate to shooting down an unmanned drone" and "we didn't have a man or woman in the drone. It would have made a big, big difference." Therefore, the Iranian leadership could reasonably foresee that the inevitable destruction of Iran's military aircraft by Israel would be relatively less likely to enrage the Iranian public and put political pressure on the government to strongly retaliate against Israel for the loss of Iranian life.

For all of these reasons and others, research shows that drones are relatively low on the escalation ladder compared to ground attacks or strikes from inhabited aircraft. The use of drones, along with the Iranian government's declaration following the strike that "the matter [with Israel] can be deemed concluded," helps serve Iran's broader strategic goal of limiting escalation, even if the attack was ineffective from a military perspective.

Iran's attack might also further another strategic political goal—maintaining a strong reputation for resolve. Many leaders strive to foster a reputation for strength for themselves and their countries by using military force, believing (even if mistakenly) that doing so can help deter foreign aggression. Following the Israeli military strike in Syria that killed two high-level Iranian military commanders, Iranian leadership may have believed doing nothing would harm Iran's image and be perceived as backing down. While impotent militarily, Iran's attack may have helped achieve this goal by demonstrating its willingness to "do something." As Iran expert Nicole Grajewski said, Iran's attack appears to have been "more concerned about symbolism than military destruction."

The Lessons of Iran's Attack for Modern Warfare

In sum, despite the meager military impact of Iran's strike, it may yet serve Iran's broader political goals. But much depends on whether Israel is willing to avoid taking additional actions that might cause the conflict to escalate into a wider regional war. The impact on Iran's reputation is also contingent on how the international community perceives Iran's initial attack and response—or lack thereof—to Israel's counterattack. While the unprecedented nature of the original Iranian attack on Israeli territory could bolster the country's reputation for resolve, Iran's transparent attempts at escalation management could undermine it. The fecklessness of Iran's attack could also end up harming its reputation for military effectiveness and thus undercut the credibility of its future threats.

In any case, the most interesting aspect of the attack may be what it portends for the future of warfare. The alleged offensive advantage current-generation drones provide over the defense is overrated, but a new era where drones can operate autonomously in coordinated large-scale swarms is

coming. To keep pace, defenders will need to continue to innovate cost-effective counter-drone technologies, including the possibility of using drones directly to destroy other drones. Sporadic drone-on-drone "dogfights" have already occurred in the Russia-Ukraine War and may offer a preview of the next generation of remote warfare.

Despite the *military* deficiencies of contemporary drones, their *political* utility will continue to be a defining element of modern warfare and statecraft well into the future. As Jacquelyn Schneider said, "These systems exist not because they are invincible, but instead because they decrease political risk for decision makers." By reducing the financial and human costs of conflict, increasing public support for the use of force, and lessening the chances of escalation, drones are having a transformational effect on international politics.

Winning the Tactical Reconnaissance-Strike Fight: Lessons from Centaur Squadron

George Pavlakis and Randall Towles

Picture kilometer-long columns of destroyed tanks and infantry fighting vehicles. Drones fly overhead while electromagnetic sensors silently parse through frantic radio transmissions. Thousands of soldiers are massed for an attack, only to stall under pummeling indirect fires. This scene could easily describe contemporary combat as warfare's changing character makes reconnaissance and strike platforms available to any potential US adversary. But rather than an anecdote from a distant conflict, this scenario is what the 11th Armored Cavalry Regiment "Blackhorse"—the National Training Center's (NTC) resident opposing force unit—has begun to inflict on rotational training units (RTUs). At NTC, the realities of reconnaissance-strike battle are painfully present, posing a challenge for RTUs that can prepare them to face the real threat on future battlefields.

Centaur Squadron, Blackhorse's purpose-built reconnaissance-strike complex, organically combines wheeled antitank and armored transport vehicles, scouts, unmanned aircraft system (UAS) operators, and electronic warfare (EW) assets. These platforms offer a combination of high tactical mobility, long-range observation, and dense firepower that feeds directly into the regimental targeting and integration cell to complete the kill chain. Centaur can also expand depending on mission variables to include light infantry, mortar carriers, and engineers.

During NTC decisive action rotation 25-07 in May 2025, we experienced Centaur's power firsthand while attached to the squadron. Fighting against a US Army armored brigade combat team (ABCT), our own organic formation, the connection between reconnaissance-strike battle theory and lethal battlefield effects quickly became apparent. Just as NTC has adapted to replicate the emerging battlefield's technological and organizational realities, ABCTs—and other brigade combat teams (BCTs)—can leverage emerging technologies too. Multifunctional reconnaissance-strike companies, combining mobile infantry, reconnaissance and strike UAS, and EW assets, can enable the brigade combat team to win the reconnaissance-strike battle, enabling decisive combined arms maneuver. Though only one anecdote from the field, we believe our experience fighting with Centaur Squadron holds important lessons for how BCTs can prepare to win their fight at the NTC and in large-scale combat operations.

Centaur Squadron and Reconnaissance-Strike Battle

Reconnaissance-strike battle conceptually connects multidomain operations from the strategic and operational echelons to the tactical. Where multidomain operations doctrine integrates joint capabilities across the air, space, land, sea, and information domains, reconnaissance-strike battle synchronizes and employs mission-relevant multidomain capabilities at the tactical level. In a military-technological environment where the US Army's adversaries possess the tools and organizational structures needed to create a reconnaissance-strike complex, the reconnaissance-strike battle will see friendly and opposing reconnaissance-strike complexes duel to establish multidomain superiority over one another. The side that gains multidomain superiority will gain the opportunity to exercise combined arms maneuver on the battlefield.

Centaur Squadron reflects this emerging dynamic with its organic fusion of sensors and shooters, along with its direct organizational linkage to the regimental targeting and integration cell at the kill chain's center. Centaur's five operational principles—flexible task organization, manned-unmanned teaming, layered reconnaissance, intelligence-derived maneuver, and tactical control of operational-level enablers—allow it to rapidly deploy into the division tactical group security zone, gain and maintain contact with RTU elements, and attrit them using manned and unmanned strike systems. As the RTU attempts to deploy itself, Centaur Squadron fights the reconnaissance-strike battle across its depth. Denying both the RTU's own reconnaissance efforts and its attempts to mass combat power for combined arms maneuver, Centaur forces it to culminate prematurely.

Infantry in the Reconnaissance-Strike Complex

During our rotation as guest Blackhorse light infantry, our company fought the reconnaissance-strike battle as part of Centaur Squadron. Light infantry provides two important qualities that complement other elements of the reconnaissance-strike complex: Its small signature makes it

highly survivable against multidomain threats, and it can carry multidomain sensor and strike payloads deep into the battlefield's most restricted terrain. While antitank scout vehicles are pushing deep into the division tactical group reconnaissance zone, road-mobile light infantry forces follow close behind to establish mountaintop observation posts from which they can sense and identify enemy elements miles away. Infantry forces can carry a variety of man-portable equipment with them, from first-person-view drones to antitank missiles and radio direction finders. Importantly, light infantry forces can immediately shift from fighting the broader reconnaissance-strike battle to repelling enemy attacks on their positions or taking the ground fight to the enemy. This factor, combined with infantry's inherently low signature compared to mounted units, makes it highly lethal and survivable on a transparent battlefield where armored formations have proven vulnerable to precision-strike kill chains.

Evidence of infantry's importance in the reconnaissance-strike battle is already apparent from the ongoing Russo-Ukrainian War. At a critical moment in Ukraine's 2023 summer counteroffensive, a single Russian infantry platoon emplaced on a hilltop stalled the Ukrainian Army's push south for twenty-four hours, buying enough time for the Russians to further entrench their lines behind the hilltop while continuing to attrit Ukrainian forces using strike assets. Both Ukraine and Russia rely on light infantry to progress forward on a battlefield where kill chains detect and strike armored and mechanized elements miles from the forward line of troops. Dismounted infantry formations disperse their combat power across dozens of individual soldiers in a platoon or company, can conceal and entrench themselves in restricted terrain that enhances survivability, and are still able to bring significant firepower to bear at long ranges using missiles and attack UAS.

Centaur's Light Infantry at NTC

Though Centaur Squadron's approach, and the emerging battlefield conditions facing the Army, appear daunting, BCTs have the capacity to tackle this challenge. During rotation 25-07's Phase I, Centaur Squadron's light infantry platoons seized key terrain overlooking the major east-west avenues of approach, providing up to fifteen kilometers of unobstructed visual observation. Antitank trucks then flowed through the mountain passes, penetrating even deeper into the security zone. This combination disrupted the ABCT as its units uncoiled, preventing them from initiating forward movement, let alone attacking Centaur's infantry in the passes.

Although scout elements made some demonstrations against the infantry positions on the ridgelines, at no point did the RTU make a discernible attempt to dislodge us or our sister platoons from our observation posts. This was to their detriment, as we could observe almost the entire ABCT's frontage from several kilometers away, providing advance warning of massing combat power for a concentrated push across the valley and the opportunity to pass targeting information up the kill chain.

Phase II featured the RTU's attempt to breach Blackhorse's main defensive line and capture the city of Razish. Blackhorse tanks and infantry fighting vehicles monitored their engagement areas while engineers had already emplaced antitank obstacles including dragon's teeth, mines, and ditches. Once again, Centaur Squadron emplaced its light infantry in mountainous, restricted terrain overlooking the main defensive line's approaches, as antitank trucks drove forward deep into the reconnaissance zone. Seeing RTU forces creep their way toward the breach, the antitank trucks could easily identify their disposition and objectives. Company after company of tanks and infantry fighting vehicles arrayed in neat formations may have been easier for their commanders to control and maneuver, but they were impossible to hide from Centaur Squadron.

After successfully attriting the ABCT's first elements bearing down on the breach, the antitank trucks retrograded and handed the fight off to the light infantry. From their position, the platoons could see almost a dozen kilometers across the valley approaching the breach and passed accurate fires targets up to feed the kill chain. When what survived of this element crossed within two kilometers, the infantry engaged with antitank missiles, stalling the breach. Centaur Squadron, with the tank company on the near side of the breach, was able to bottle up the RTU's main effort by feeding the Blackhorse kill chain faster than the ABCT could respond, winning the reconnaissance-strike battle. The RTU made piecemeal attempts to dislodge the infantry from the hills, but never enough to prevent them from disrupting the breach or maneuvering on rear-echelon high-value targets including air-defense Strykers. Additionally, if infantry or scout elements from the ABCT had seized our hilltops, they would have had nearly unobstructed observation of every Blackhorse battle position overlooking the main defensive line, and could have passed targeting information up their own kill chain.

This pattern repeated itself during Phase III. Blackhorse set another main defensive line with its tank and mechanized infantry battalions overlooking engagement areas in a valley while Centaur Squadron occupied key terrain and pushed antitank trucks forward. As the RTU maneuvered westward toward its breach, the massed formations of tanks and infantry fighting vehicles evoked real-world scenes of Russian armor massing for attacks on cities such as Vuhledar and being destroyed in the process.

How BCTs Can Prepare to Win the Reconnaissance-Strike Battle

Centaur Squadron combines emerging military technologies and existing platforms to enable rapid target identification as part of Blackhorse's reconnaissance-strike complex. Observation—whether through a UAS, EW collection, or an infantryman's binoculars—is Centaur's most dangerous weapon, and one that RTUs can take meaningful steps to counter by simply changing their behavior. As a BCT prepares for its fight, identifying key terrain that provides long-range unobstructed observation is critical. Centaur will almost certainly emplace observers there, who will identify and direct fires on any RTU elements within range. RTU forces can effectively use terrain to obscure themselves from these points as they approach, something the ABCT came close to doing on several occasions during rotation 25-07. In addition, as we ourselves realized while defending the main defensive line during Phase II, leaders must recognize that key terrain is as

relevant for friendly forces as it is for the enemy. If an enemy infantry platoon with tank and preparatory fires support had attacked our hilltop and seized it, it could have decimated Blackhorse's entire defense in that sector of the main defensive line by passing accurate targeting information to the RTU kill chain.

BCTs can only effectively maneuver on Centaur Squadron, however, if they eschew massed formations for smaller maneuver elements, potentially at the section level or below. Centaur Squadron's observers can easily spot a tank or mechanized infantry platoon up to ten kilometers away from an elevated observation post. Companies and battalions are visually observable from even farther away, and only the weather and platform range limit UAS or EW systems. In comparison, light infantry's small signature allows small units to infiltrate great distances with a far lower chance of detection by the enemy's reconnaissance-strike complex. The urge to consolidate forces, easing control over them, is understandable, especially when navigating long-distance movements in the desert. Mission command, however, offers a promising alternative. Dispersing individual vehicles at release points outside Centaur's sectors in the reconnaissance zone, with clear commander's intent and a rally point to mass combat power just before the attack, can help commanders maintain surprise and audacity. The Russian Army has already learned this lesson in Ukraine, often releasing entire sections or platoons to maneuver on an objective as individual soldiers to avoid omnipresent attack UAS destroying them in massed formations.

Centaur's natural habitat is restricted terrain. Whether it be hilltops, rock piles, villages, or draws, restricted terrain allows Centaur Squadron to hide in plain sight while observing the RTU at a distance. If BCT elements can see it, they can be seen from it. BCTs cannot afford to neglect restricted terrain and must at least actively reconnoiter these positions to enable their own freedom of maneuver elsewhere on the battlefield. This means that infantrymen, especially those in a reconnaissance-strike complex, must be physically fit enough to move several miles across Class 2 and 3 terrain just to reach the fight. The tank fight on open plains depends on how effectively infantry can dislodge observers nestled in hilltops, while the infantry fight in the cities depends on well the BCT can observe avenues of approach leading to them.

The RTU was eventually able to take the fight to Centaur Squadron using its own emerging technologies. Shortly after our platoon observed unknown reconnaissance drones near our positions during Phase I, drone-observed indirect fires destroyed the company's supply trucks and command post a few hundred meters away. Had the RTU attacked at this moment, it could have dislodged the infantry from the key terrain, opened the passes, and maneuvered armor-infantry teams through them unopposed. Centaur Squadron kills the enemy by observing its forces at a distance, including with UAS and EW systems, but is just as vulnerable to ABCTs who can effectively use theirs.

Winning The First Fight

Several BCTs across the Army have already made progress toward winning the reconnaissance-strike battle. We believe that the multifunctional reconnaissance-strike company builds on these existing efforts while reflecting the payload-agnostic nature of small UAS platforms. As a brigade-level asset, this company would combine the reduced signature and high mobility of infantry with the survivability and firepower of a combat formation armed with antitank missiles and strike UAS. The soldiers in this formation must be physically fit enough to traverse miles of difficult terrain and reach observation posts close to the enemy's manned zone, from where they can feed targeting information up the kill chain while engaging with missiles and strike UAS. They must also be competent and well-trained enough to accomplish this task, exercising mission command with a high degree of autonomy. A unit with the right people, equipment, and training can win the reconnaissance-strike duel, enabling combined arms maneuver for the rest of the BCT. Thanks to recent policy changes aimed at "unleashing U.S. military drone dominance," brigade commanders can foster this change at their level and equip their formations for reconnaissance-strike battle.

NTC rotations have tested generations of Army leaders in the closest thing possible to real ground combat at the brigade scale. The next generation of Army leaders, preparing to fight a war in which local American military supremacy is not guaranteed, can expect challenging and technologically realistic training when they come to NTC. Centaur Squadron represents Blackhorse's nod to emerging military technologies and the novel task organizations that maximize their effectiveness on the battlefield. This poses a major challenge to any BCT, but its value to them is as a forcing function to adapt or fail. Learning to fight Centaur in this training environment, a feat within reach of any brigade in the Army, will prepare units to counter dangerous emerging threats on a rapidly changing battlefield.

Imagining a US Army Drone Corps

Joshua Suthoff

In February 2024, Ukrainian President Volodymyr Zelenskyy announced the creation of the Unmanned Systems Forces. It is no secret that the Ukrainian military has used drones to great effect. Its units continue to innovate with drone tactics, techniques, and procedures and effects in the air, land, and maritime domains. Both belligerents in the Russia-Ukraine War have pledged to build over a million aerial drones each year to fill the skies. Even with the extremely innovative use of the drones (mine laying, incendiary delivery) already observed in Ukraine, history will show that the most important attribute of drones has been their ability to serve as economy-of-force systems. In a grinding war of attrition, drones have allowed the Ukrainian military to protect its limited combat power and threaten a much larger combat force across multiple domains.

The Unmanned Systems Forces that Zelenskyy announced amount, effectively, to a drone corps. US policymakers have taken note of the effectiveness of drones in the conflict and a drone corps may also be coming to the US Army. A draft 2025 National Defense Authorization Act (NDAA) from the US Congress directed the Army to establish such a corps as a basic branch. The language did not make the US Senate's version of the NDAA—however, it is an understatement to say that drones and their effects are here to stay on the battlefield. Drones may not be revolutionary in their impacts, but they are a creditable and enduring enabler—an enabler that continues to threaten the hegemony of traditional branches like artillery and the once dominant mass of infantry and armor. The Army and DoD more broadly are striving to innovate with drones, but standardization, training, and tactics vary. The DoD Replicator program's primary mission is to increase the available drone inventory, but how effective is a deep magazine of drones without subject matter experts to operate and employ? This trend of drones' growing impact will likely continue as they are paired with AI and other technologies. But in an imagined near future where the NDAA has passed with the drone corps language intact, what form should such a corps take?

The Needs Statement

It is no secret that the US Army (and other services) faces recruiting challenges. The force would certainly struggle to quickly fill and train its ranks for a mass mobilization. This problem is compounded by the increasing speed with which the character of war is changing—which means the Army would not just need a sufficient number of people in the event of a large-scale conflict, but people whose knowledge grows at a pace commensurate with battlefield evolution. A case in point is the war in Ukraine's repeated demonstration of actions and counteractions as both sides seek to ensure drone primacy. The most recent example is fiber-optic tethered drones largely immune to electronic warfare effects.

Drones, if managed appropriately, are a quick, cheap, economy-of-force capability, but one evolving rapidly enough that it requires a focused and professional stewardship—a drone corps. As the Army closes most of its cavalry squadrons there also remains the requirement for reconnaissance and security operations. These are historically an economy-of-force mission and drones are well suited for it when supported with traditional enablers. In a cavalry squadron, the planning, collection, and dissemination of the intelligence requirements was a primary focus of the squadron staff. Now the task is spread across multiple headquarters with no real unified ownership. A drone corps would help to fill this gap.

The creation of a drone corps is the next evolutionary step in the critical management of these warfighting systems. The training, maintenance, and implementation cannot be the secondary duty of a brigade aviation officer or a proactive noncommissioned officer or officer in a battalion. A drone corps will ensure development of systems and branch professionals along avenues across the DOTMLPF spectrum (doctrine, organization, training, materiel, leadership and education, personnel, and facilities). Infantry, sustainment, and fires battalions are all led by professional officers and noncommissioned officers from those branches. Drones and their employment have evolved to the point where those domain systems need the same stewardship. The Army cannot afford a disjointed implementation of drone doctrine or of tactics, techniques, and procedures. The systems need a champion to drive progress inside the Army and refine requirements for industry.

Waypoints to Follow

So, what would a professional drone corps look like? Among US special operations forces (SOF), the SOF truths provide a critical baseline understanding of purpose and a north star for the special operations community's culture. A drone corps would be well served by replicating this model and building itself around a fundamental set of drone truths.

1. *Drone warfare will be an enduring capability and threat on the future battlefield.* Drones are cost-efficient, simple, and quickly mass-produced. They allow individuals and states the ability to compete with larger adversaries. People will find a way to keep drones in the fight.

2. *Drones do not replace the warfighter.* Even with drones on the battlefield, manned maneuver forces will still be needed to retain ground and fight. Humans must still be on the loop for maximum efficiency. AI is efficient, but—because it does not precisely replicate the cognition of humans—insufficient.

3. *Drones are a powerful economy-of-force capability and constant enabler.* Drones don't need to rest. Instead, they allow manned forces to rest, refit, and transition to the next fight. A cheap drone can always be sent forward, preserving combat power and saving lives.

4. *Drone professionals, effects, and synchronization are not created instantly, but require training, planning, and time.* Drone systems are relatively cheap and it does not take long to train an operator. However, it does take time to achieve synchronization and effects. Additional-duty operators found in maneuver battalions vary in quality, flight time, and capability. Regardless of operators' skill, they will eventually leave their units. Subject-matter experts and repetition are required to ensure drones are well positioned and the data they provide is appropriately analyzed.

If these truths form the basis on which to establish a drone corps, that corps should also be constructed with several foundational principles in mind. First, the drone branch is responsible for the operations, training, and testing of the systems. This includes air, ground, offensive, defensive, unattended sensor, and resupply drones.

Second, the drone branch is fundamentally a maneuver branch with a focus on targeting and synchronization. It does not replace the existing maneuver branches, but works to enable and support their traditional offensive and defensive characteristics. The lethality and agility of drones (especially in mass) dictate that drones move from the intelligence warfighting function to maneuver. It made sense only a few years ago to align drones under military intelligence when they were primarily intelligence, surveillance, and reconnaissance platforms. Drones in their current multirole configurations, however, are more comparable in function to a main battle tank moving forward, engaging targets, positioning wingmen, and developing the situation. Sensing is now potentially only one aspect of a drone. Maneuvering air and ground drones and planning their effects follows along the same lines as doing so with infantry and maneuver formations. Drones require a maneuver mindset.

Third, recruiting for the branch should focus initially on maneuver and intelligence officers and noncommissioned officers. In the longer term, public awareness of—and interest in—drones means that the creation of a drone corps could increase overall recruiting and retention.

Fourth, the drone corps operates within maneuver formations or in drone-specific units.

And fifth, the drone branch is not a subset of the aviation branch, but works in concert with it to deconflict and simplify airspace management. Drones now operate across all domains, well beyond the air littoral where they were once almost exclusively found.

A Tactical Formation, Not a Niche Specialty

There are ways of meeting the letter of the draft NDAA language while stopping short of establishing a drone corps—simply creating the branch, for example, by changing the military occupational specialty of a few soldiers in maneuver formations. This would be similar to some of the Army's lower-density military occupational specialties, like electronic warfare. The easiest path would be to rebranch soldiers in formations like the human-machine integration platoons. Individual subject-matter experts might be gathered on select headquarters staffs where other small-density specialties are consolidated. Maneuver companies would maintain organic drones to provide local situational awareness and strike capabilities. These operators could be either soldiers designated by the new drone military occupational specialty or additional-duty operators. Focusing drones only at the maneuver battalions is limiting the true potential of drones.

This simple and straightforward course of action, however, would not help with the imperatives of standardization and training. Moreover, the Army would be missing an opportunity to build a small lethal, formation that can enable other maneuver units or operate itself as a unit of action. Expanding task organization for drone use is the true spirit of the NDAA's draft language.

As an Army we are at a critical inflection point and have an opportunity to build a lethal enabling force. A more expansive course of action would involve creating drone units that can operate independently or augment brigade formations to fully leverage the situational awareness and strike capability of the systems. In a zero-growth environment with no major budgetary reallocations, the ready solution is the consolidation of the human-machine integration platoons across a division to build a robotics recon strike squadron (R2S2). An additional manning solution could incorporate the intelligence, surveillance, and reconnaissance company in the intelligence and electronic warfare battalion. An R2S2 would be small formation task-organized under a division headquarters, commanded by a lieutenant colonel.

There would be two primary missions of the R2S2. The first is to standardize training, maintenance, and equipping across a division, ensuring that the additional-duty operators within maneuver formations are proficient. This would be similar to a DIVARTY (division artillery) concept, and in fact the fires enterprise more broadly is an excellent model for a drone crops. Fires professionals in maneuver battalions, brigades, and separate fires battalions ensure that the warfighting function is appropriately trained and leveraged.

The second mission is to deploy human-machine integration platoons or companies in front of or supporting maneuver battalions to conduct targeting and reconnaissance and surveillance operations for supported brigades or divisions. This provides an economy-of-force element that creates decision space for commanders at echelon.

An R2S2 assigned to a division headquarters would be uniquely positioned to sense with robots, coordinate the intelligence picture with the adjacent intelligence and electronic warfare battalion, and strike with division, corps, and joint assets. More importantly, establishing a headquarters to manage the data provided by drones in all domains ensures that critical data can be translated into intelligence requirements or strikes. This headquarters would close a

critical gap created with the deactivation of traditional cavalry squadrons. US Navy senior leaders in the Indo-Pacific region have advanced a vision, which they describe as a "hellscape" and which acknowledges drones' rapid economy-of-force potential to their ability to provide time and space for the joint force commander in the event of a Chinese incursion into Taiwan. An R2S2 would provide the Army with a deployable land- and air-based "hellscape" option at the next conflict zone.

The R2S2 does not need to be a large formation and should leverage other sustainment functions within division headquarters formations to reduce overall manning requirements. For example, the R2S2 could receive routine S-1 or S-4 support from the established headquarters battalion, intelligence and electronic warfare battalion, or DIVARTY. The R2S2 headquarters should be designed to focus primarily on sensing and striking for supported divisions or brigades. Human-machine integration platoons and companies within the RS2 would be equipped with infantry squad vehicles, multiple drones, and secure but unclassified communications systems to allow them to rapidly integrate and support maneuver battalions. A light package would ensure the platoons and the R2S2 can reposition or deploy quickly in support of economy-of-force missions and support to lodgments. Besides sense-and-strike missions, an R2S2 could also support mine emplacement and critical resupply operations for a division headquarters.

A third course of action would involve consolidating both human-machine integration platoons and battalion scouts under an R2S2 headquarters. This would further increase the reconnaissance and strike capability of the unit and ensure brigades and divisions have the best cued and mixed intelligence picture. It would provide a true all-weather sense-and-strike formation.

Completing the Kill Web

Establishment of a drone corps and R2S2 units would provide a key link in the Army's evolving kill web. A well-positioned R2S2 with drones forward in all domains would establish a clear threat picture for a higher headquarters. While providing immediate kinetic effects with drone munitions, it would also provide strike recommendations and sensing for key shaping formations like multi-domain tasks forces and rocket battalions. Establishing subject-matter experts and utilizing the fire support model reduces cognitive load and training requirements for maneuver battalions. The training, maintenance, and use of drones within maneuver battalions are absolutely critical, but these activities come at the cost of time and distract from those battalions' mission-essential tasks. Drone professionals and a focused R2S2 headquarters would ensure drones and operators are ready when needed. Having an R2S2 look forward while a brigade or division deploys or transitions would ensure that tempo is maintained. One can find multiple examples of Russian armored and infantry attacks being disrupted and destroyed by waves of drones with minimal or no friendly maneuver support. Pairing the use of drones with a combined arms mentality of the US Army increases overall lethality and maximizes drone effects. An R2S2 would also provide a continuous layered and joint defense, as well as a sense-and-strike capability for a division.

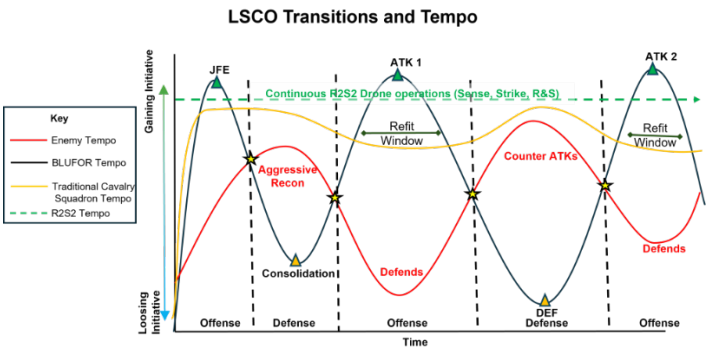

In short, the R2S2 would provide a small drone–focused package to tie together the Army's kill web. Small and agile units would move and hide on the battlefield while pushing drones forward. Equipped with one-way drones and linked to the fires network, an R2S2 would punch above its weight and buy time for a higher headquarters.

The Next Evolutional Step

New weapons and platforms routinely drive change in the character of warfare. The introduction of tanks and planes came with great consternation as to how they should be used. The incorporation of drones is no different: they are a powerful asset that must be managed by dedicated professionals, similar to fires and aviation branches. The drone corps, even if established with a built-in R2S2 concept, is not the final step, but the first in a continual process of improvement and transformation. In a resource- and personnel-constrained environment, the Army cannot afford to mismanage its drone force. It must fundamentally reconceptualize what it considers maneuver branches, and how it organizes and fights maneuver formations. Drones will remain a major feature of the near-future battlefield. To be prepared for that battlefield, the Army needs a drone corps and its subject-matter expertise.

Understanding the Counterdrone Fight: Insights from Combat in Iraq and Syria

D. Max Ferguson and Russell Lemler

Between August 2023 and April 2024, 2nd Brigade Combat Team, 10th Mountain Division was deployed across Iraq and Syria in support of Operation Inherent Resolve. During that time, state-sponsored militia groups launched over 170 attacks against a small network of coalition bases that 2/10 was responsible for defending. The brigade's deployment represents the most recent and direct experience of any US Army unit defending against drone attacks and, consequently, an important set of lessons on countering and defending against rockets, missiles, and drones of all sizes.

The soldiers of 2/10 experienced a wide range of enemy attacks against conventional munitions, from rockets and mortars to cluster munitions and short-range ballistic missiles. But the enemy's weapon of choice was the one-way attack unmanned aircraft system (OWAUAS). These drones were mostly little, propeller-driven, fixed-wing craft made of carbon fiber, metal, and plastic. They flew low, sometimes less than a hundred feet off the ground, and depending on the type, their wingspan was a few feet to a few meters. Their US military equivalents are between a Scan Eagle and a Shadow. These systems have no landing gear because they're designed to land on their noses with a bang.

UAS Chart

	GROUP 1	GROUP 2	GROUP 3	GROUP 4	GROUP 5
Range	10-30 km	100-700 km	1,000-2,500 km	200-400 km	>2,000 km
Speed	< 100 knots	< 250 knots	< 250 knots	> 250 knots	> 250 knots
Altitude	< 1,200 ft.	< 3,500 ft.	< 18,000 ft MSL	< 18,000 ft MSL	> 18,000 ft.
Endurance	1-2 hours	18 hours	Approximately 11.5 hours	12 hours	14 hours
Payload	2 KG	5 KG	30-50 KG	100-300 KG	300+ KG
US Equivalent	DJI SBS Raven	Switch Blade Scan Eagle	Shadow RQ-7B	Grey Eagle MQ-1C	Reaper MQ-9

Figure 1. Group 1–5 Unmanned Aircraft System Overview

The low-flying, low-cost, highly accurate, and prolific drones are irresistibly effective. Despite their small size, one-way attack drones on the battlefield today have tremendous range. Small to mid-sized one-way attack OWAUAS can travel as far as 2,500 kilometers, distances more akin to land attack cruise missiles and ballistic missiles than any tube-based artillery. Their versatility, reach, cost, and precision will increasingly make them appealing options for any modern combatant, no matter its global stature or military size.

Drawing out the Lessons for Large-Scale Combat Operations

In one sense, 2/10 experienced an unprecedented number of air attacks. In a few short months, 2/10 accumulated more combat experience in defending against OWAUAS than any unit in the Army. Yet this experience may soon seem miniscule compared to future conflicts. The number of attacks 2/10 handled over four months could conceivably occur in just a few days—or less. In Ukraine, for example, Russia has launched dozens of drones in single strikes. The US Army should prepare to encounter this frequency and scope of attacks in a future large-scale combat operations (LSCO) fight.

Context matters. The soldiers of 2/10 were employed in a very distinct role, defending from established fixed sites with reliable connectivity, hardened bunkers, and air mobility that was limited by electromagnetic interference, adversary surface-to-air weapons, and political constraints but still only partially contested. Most importantly, 2/10's fight was a purely defensive one. Even when some of the counter–unmanned aircraft systems (C-UAS) capabilities were designed to be mobile, they were employed as part of a static defense. The base defense operations centers were well established and operated in controlled environments.

The attacks were continual but measured, nothing like the artillery barrages that Russia employs in Ukraine or what we know is possible for drone swarms from a well-supplied enemy. The OWAUAS salvos against 2/10 came mostly in singles and doubles. Rockets and missiles came in sporadic batches, sometimes a little over a dozen at a time. The consistent but low volume attacks were by the enemy's design under the unique circumstances of the conflict. Thomas Friedman described the conflict as a "shadow war" and "the

most dangerous game of chicken going on anywhere on the planet today" after visiting several CJTF-OIR (Combined Joint Task Force–Operation Inherent Resolve) outstations with the commander of US Central Command in February 2024. Escalation was carefully managed on all sides to apply pressure as part of a campaign of coercion. The types of attacks that 2/10 experienced were distinct to this particular conflict, which itself does not represent what the US Army will face during LSCO.

Still, we must pull lessons where we can, not least because artillery and aerial attacks will dominate in the next fight. Sifting through the unique characteristics of the C-UAS fight that 2/10 experienced in OIR reveals fundamental characteristics about C-UAS operations that need to be captured and applied to US Army operations.

The Three Types of Defenses: Kinetic, Nonkinetic, and Shelter

At the end of the day, there are three ways to defend against one-way attack drones: you can shoot them down, you can hit them with electronic interference, or you can seek shelter and absorb the hits.

Shooting a drone down using kinetic defeat comes in the form of missiles, guns, and directed energy lasers. Nonkinetic electronic defeat options vary depending on how the drone is controlled. The effects vary as well. If the drone is actively piloted, you can target the link between the ground control station and the aircraft. If the drone is GPS-guided, you can target the link between the GPS and the satellite. The enemy is constantly adapting technology and probing vulnerabilities in sensors and radar coverage. In just nine months, 2/10 saw advances in attack vector choice, GPS hardening, and other adaptations.

Seeking shelter varies by type of operation. In a deliberate defense, it means going to reinforced bunkers. On the offense or a hasty defense, shelter might vary from a rapidly dug foxhole to occupying a hard clad building or commandeering a basement or cellar.

All of these methods of defense rely first and foremost on detection. Bunkers only work if you're inside one. You can only alert personnel to occupy a bunker if your sensors detect the threat with ample time and your alerts systems are working properly. The earlier you detect a threat (drone, rocket, missile, or artillery), the sooner you can alert the force to seek shelter while the air defense operators work to employ their systems to defeat the threat.

Overall, like any other battle plan, C-UAS operations are a multilayered multifaceted defense in depth. The best way to protect the force is through a combination of both active and passive defensive measures. Below are some observations on these different measures for leaders to consider.

Kinetic Defeat Considerations

Each kinetic option has limitations and constraints, just like any other weapon system. Missiles are most effective at range and provide standoff distance, but different systems have different time requirements to engage a threat. Patriot batteries have the best range and altitude but are designed for ballistic missiles and defense against high-performance aircraft (jets and Group 4–5 drones).

There are smaller missiles designed specifically for low- to mid-altitude drones (Group 1–3) and low-performance, principally propeller-driven aircraft. These drone interceptors range beyond ten kilometers with the right terrain and radar coverage. Initial missile variants took an uncomfortable amount of time to spin up before being ready to launch. But when these missiles do find their targets, they produce impressive displays of aerial acrobatics and precision.

These short-range air defense missiles were lifesavers for 2/10's C-UAS fight, but they have important limitations in a LSCO context. Producing these missiles is a sophisticated and labor-intensive effort. Current missile variants can cost upwards of $100,000 each, about ten times the cost of each drone they defeat. Manufacturing outputs are sufficient for the consumption rates in OIR, but a future conflict could see an exponential demand for such systems on a global scale.

Guns, like the C-RAM (Counter Rocket, Artillery, Mortar), are more responsive and cost efficient than missiles, but their range is uncomfortably close. The Land Phalanx Weapon System (LPWS) fires 20-millimeter self-detonating rounds from a six-barrel Gatling gun. The LPWS is a variant of the Navy's Phalanx Close-In Weapon System originally designed in the 1970s. As the name implies, it is only capable of defending against immediate threats (roughly within one kilometer). The LPWS operates on a trailer, with a limited mobile application, and is best used to defend key nodes rather than mobile troop formations at the edge of battle.

Directed-energy weapons are steadily emerging on the battlefield to complement traditional kinetic defeat options. During its deployment, 2/10 tested and employed several early variants of these short-range, directed-energy systems as both fixed-site and mobile platforms. They are futuristic and impressive in concept but work more as a slow burn than a Star Wars weapon. A laser might knock a drone out of the sky or deflect it enough to miss the intended target, but it might also just singe the paint off the drone before it dives into its destination. The drone's distance, speed, and construction material, combined with the laser's power output, time of acquisition, and target-lock capability, will determine effectiveness.

The US Army will face constraints with the employment of lasers in a LSCO environment. The amount of available power is important, so austere sites and mobile systems for land applications will require more technological advancements. Moreover, directed-energy weapons are hard to ruggedize and rely on sensitive components to concentrate the energy for effects.

Therefore, lasers have a promising future in the Army's C-UAS arsenal, but they are still in the early stages of development and fielding. Their dispersion across the battlefield will depend on advancements in mobile power generation and whether the systems can be built to withstand the rough-and-tumble nature of combined arms maneuver.

Nonkinetic Defeat Considerations

There are several benefits to nonkinetic electronic warfare (EW) systems in the C-UAS fight, but these also have several important limitations and nuances. Electronic defeat equipment works by disrupting the link between the drone and the control station or by interfering with the drone's navigation. Operationally, 2/10 found these systems most effective against smaller drones, especially commercially developed quadcopters and reconnaissance drones controlled by a ground station. When OWAUAS are guided by preprogrammed waypoints, some electronic countermeasures might still work, but kinetic defeat options are currently more effective.

A key consideration that 2/10 faced when employing various forms of EW was the secondary effect some electronic weapons can have on friendly communication systems. Some EW platforms use radar detections to target specific drones or interrupt select frequencies to redirect or disable drones outright. If these precision options are not available or not effective, other EW countermeasures include broad electromagnetic interference capabilities. Like any other electronic countermeasure system—for example, Duke systems designed to defend against improvised explosive devices—C-UAS EW platforms require regular updates and new fills to stay up with the regional threats.

Note that while more powerful EW measures can achieve the desired results, the effects sometimes come at a cost. Emitting a powerful electromagnetic interference (EMI) burst can be like chemotherapy for the radio waves. Some systems can zap everything in range and may just as readily interfere with friendly electronic and communication devices as with enemy signals. If the base or site's sense-and-warn systems rely on wireless signal traffic, an electromagnetic interference burst may sever the signal to alert friendly forces of the incoming threat.

As these electronic countermeasures continue to transform, their effectiveness will ebb and flow, and opponents will evolve their shielding against electromagnetic interference. The frequencies and terminal guidance methods will change. It will be no different than the twenty years of improvised explosive device adaptations that the US Army defended against in Iraq and Afghanistan. It will be the same game, but with the threat now coming from the sky, not the ground.

Seeking Shelter: A Matter of Bunkers and Warnings

Despite all the expensive technology involved in defeating drone threats, some of the simplest and cheapest options—like sandbags and concrete—are still among the best protection measures. Bunkers remain an essential component to static defenses. The men and women of 2/10 can attest that as crude as they are, they work—as long as people were inside them.

The members of OIR spent considerable time in bunkers during 2/10's deployment—sometimes multiple times a day or for hours at a time, depending on the threat. Seeking shelter became instinctive when the initial static was heard from base speakers. Soldiers and civilians came to accept the process, because these bunkers were 100 percent effective at saving lives, including instances where bunkers took direct hits from OWAUASs.

The combination of bunkers and radars were sometimes as important as concrete T-walls to 2/10 in the C-UAS fight. T-walls contain an unexpected blast from spreading and interfere with flight paths but offer no overhead coverage. Overhead protection, including reinforced bunkers and pre-detonation roofing on select buildings, were extremely important as attacks became more precise and targeted toward high-occupancy areas and key command nodes.

Bunkers came in various styles and designs, but most were concrete "C-channels" averaging six feet high and eight inches thick.

On their own, these three-sided structures are insufficient against most OWAUAS warheads and 107-millimeter rockets. Sandbag parties became the norm to reinforce C-channel bunkers with layers of sandbags or HESCO barriers. When available, the end caps were protected with welded steel doors or surrounded by mid-height concrete barriers. If you can see out of the bunker, the shrapnel can find its way in.

A good radar network tied to an effective sense-and-warn system provides advanced warning of a threat for personnel to seek shelter. On fixed sites, sense-and-warn systems can include speaker towers and indoor speaker boxes; audible coverage should be sufficient to alert someone in the shower or wake up sleeping personnel in all occupied zones of the base. In an offensive maneuver scenario, radio calls and mobile speaker boxes may need to suffice, but the expectation for rapid notification across the formation must remain the same.

Notification times will vary based on the threat. Some of the larger, fixed-wing drones were detected several minutes out, allowing sufficient time to seek shelter. Some threats get masked by terrain or hidden within a congested air picture (especially by airfields). In such instances, incoming drones were detected just seconds away from the base, providing time only to brace for immediate impact. Quadcopters and other small drones may be launched near a base to close in on a target in seconds. It may seem minor but the specific terms used by the system to communicate the threat are important and must clearly and immediately indicate what action personnel must take—whether to brace for immediate impact or seek shelter, for instance.

Notification systems provide other uses. "Big voice" systems help transmit guidance to sheltered personnel who lack any other means of communication. Announcements can instruct personnel to remain in bunkers, prepare to defend against a ground attack, deploy select response personnel, or begin unit accountability protocols. Like a fire drill in school, people go to shelter as they are, at whatever moment the alert is triggered. They might not have communication devices. They might not have pants. So sense-and-warn systems become an essential command-and-control feature in C-UAS fights.

The Limitations of Dismounted Systems

There are several dismounted systems available for both kinetic and nonkinetic defeat. These can fill critical gaps for the close-combat force but there are some important limitations to recognize about deploying these systems to defeat UAS threats.

Dismounted kinetic options include man-portable air defense systems (MANPADS) such as shoulder-fired Stinger missiles. It is important to weigh whether a dismounted Stinger team has the right training, equipment, and situational awareness to be effective against a drone attack. For example, does the MANPAD includes a day and a night sight? Soldiers may improvise by zip-tying a thermal optic to the Stinger, but that is a hasty measure at best. It is also important to understand how portable missile systems acquire and lock onto targets. Stingers are heat-seeking missiles designed to destroy helicopters and airplanes. Smaller drones may not produce the requisite heat signature to acquire "tone" with a Stinger. Also, it is important to appreciate the different variants of surface-to-air missiles. Common Stinger variants are point-detonation weapons while others are air-burst munitions. It is less likely that a point-detonated Stinger will intercept a small, low-heat-producing drone.

Nonkinetic choices available today produce directed electronic interference and defeat options with portable devices that often give off strong Starship Trooper vibes. Current "shoulder-fired" systems are useful for intercepting small quadcopters loitering overhead and conducting observation but dismounted nonkinetic C-UAS systems have limited effectiveness in defeating fast-approaching attack drones.

There are several risks to deploying dismounted teams to defeat drones. If these soldiers are not integrated with a radar network and are reliant on visual target acquisition, the risk of hitting a friendly aircraft by mistake increases. This is especially high when operating near airfields or friendly air corridors. Ideally, C-UAS systems will advance to the point where dismounted troops at the front edge of battle can employ MANPADS as part of an integrated radar network that covers the friendly force. Until then, leaders should be clear-eyed about what they may achieve (or risk) by deploying Stinger teams against OWAUAS threats. Having dismounted teams on the berm, rather than in bunkers, may feel prudent without actually being effective. Commanders must carefully consider the characteristics of the handheld equipment available to ensure it matches the anticipated drone threats.

The Engagement Process

Employing these air and missile defense systems follows all the same fundamentals of any traditional defense in depth, just with a few added variables. Success depends on a unit's ability to find and characterize threats with sufficient time to employ a hodgepodge of layered defense systems. That activity is best understood through the C-UAS process of detect, identify, decide, and defeat.

Redundant and overlapping radar systems detect air tracks as far out as possible. Battle captains then establish positive identification (typically through digital means, as most tracks populate beyond the visual range) and decide on one or more kinetic or nonkinetic means of defeating the threat.

Figure 2. The Counter-UAS Process

In the current base-centric OIR mission, this process is conducted solely by the base defense operations center (BDOC), defined in joint doctrine as "a [command-and-control] facility established by the base commander as the focal point for [force protection], security, and defense within the base boundary." The BDOCs operated by 2/10 at multiple bases throughout Syria and Iraq proved the value of having established crews in hardened command posts with

connectivity and access to multiple intelligence, surveillance, and reconnaissance platforms. The BDOC construct is effective at employing the C-UAS process; static and redundant radar systems enable early detection and positive identification of threats, rapid decision about engagement platform, and successful defeat by missile, LPWS, laser, or EW. In a LSCO fight, ground force commanders must maximize their ability to not only command and control their formations, but also execute this ever-changing air defense battle drill with varying systems.

The transitions from base commander to maneuver commander, from base boundary to area of operations, and from fixed-site command post to mobile command post are worth contemplating. The Army is fielding several movable radar and interceptor systems and actively developing new systems to be tested soon. Regardless of what equipment shows up on the hand receipt, we would be wise to understand, implement, and test to provide feedback and tailor the tools to our mission end states.

The personnel and equipment requirements for the C-UAS battle drill are in direct conflict with the understandable pressure on ground force commanders to reduce the size and electromagnetic signatures of their command posts. Even a reduced battalion or brigade command post typically encompasses multiple vehicles, generators, antennae, and tentage or shelters. At a minimum, the C-UAS fight means additional data entering the existing command post, and an added layer of decision requirements and stress.

Dispersion complicates the issue further; much of 2/10's C-UAS success was a direct result of colocating a team of nine or more BDOC personnel to monitor multiple radar systems, clear air, communicate across a base (and report upward), and conduct emergency response. A unit in LSCO may not have the luxury of colocating such teams.

Observations from Transforming in Contact

Counter-UAS systems will continue to transform and adapt for years to come. The soldiers of 2/10 lived through the experience of transforming in contact and gained some important insights worth sharing. The vignettes below describe specific issues that may not be precisely replicated in future operational environments, but they are shared to highlight some of the subtle challenges that transforming in contact creates.

Testing New Equipment on the Real McCoy

Prototype equipment was regularly tested by 2/10 during the deployment, including directed-energy lasers and various C-UAS missile systems. Live-fire tests involved model Group 1–3 UAS test drones. The bodies of these drones were not made of the same materials as enemy drones. Live-fire experiments gave a false sense of confidence in the effectiveness of the prototype weapons because these friendly UAS were easier to shoot down than the enemy drones. It is important that validation test fires include replicas of enemy systems to ensure that the prototype equipment has the desired effects.

Bearing, Altitude, Range, and Speed

The contact report in a C-UAS fight is called "BARS"—referencing the bearing, altitude, range, and speed of a suspected track. Radar systems often provide this data well before a unit can gain visual identification with current camera systems. The combination of these four elements helps leaders and operators distinguish between birds, balloons, OWAUAS, and friendly aircraft. Sometimes the altitude is the critical factor (only certain airframes can fly at thirty thousand feet) or the speed is too slow to be a fixed-wing drone. Range helps leaders understand time available to react. BARS needs to become a familiar concept the same way SALT and SALUTE reports became common knowledge across the Army.

Fourteen Clicks

Early versions of the radar software included cumbersome interfaces. Radar operators found themselves needing to perform up to fourteen clicks of a mouse to interrogate suspected tracks and deploy countermeasures during an engagement sequence. Such a laborious interface would create a deadly flaw against a swarm scenario. As operators identify these sorts of inadvertent inefficiencies built into the software, development engineers need to be informed to quickly update programs.

Knowing the Default Settings on Systems

For several months, close engagements ended prematurely with missiles terminating before hitting targets because OWAUAS were getting too near to the base. C-UAS missiles were self-detonating or refusing to launch altogether based on default engagement settings programmed into the software. Such well-intended safety features created by software engineers were not well communicated to commanders in the early stages of the conflict because the issue had never arisen before. Once the settings were explained and understood, commanders were able to adjust them on a case-by-case basis based on the commander's risk considerations. This greatly improved the kill ratio of the C-UAS engagements afterward. As new systems are fielded, programmers and operators need to know which software features are fixed and which are adjustable. Theater commanders should have the knowledge, authorities, and technicians available to adjust settings as needed. As enemy tactics and weapons evolve, commanders need expedited flexibility to make the risk-reward calculus for system modifications.

System Performance Analysis

Program engineers from the Joint Analysis Team in Huntsville, Alabama routinely provided critical after-action feedback on the performance of missiles and radars during operations. These engineers could determine the effectiveness of radar placements and diagnose issues when missiles failed to hit targets. Though the quality of their work was first class, they were far from the battlefield and access to their teams was at times delayed. The Army needs to ensure support teams are sufficiently staffed and accessible to meet the demands of battlefield commanders. If that means deploying engineers and software design teams forward, the Army should do so, especially for novel systems still undergoing field testing and evaluation. This is literal rocket science after all, and the fight evolves at combat speed. The whole team needs to be together to stay ahead of the enemy.

An Overreliance on Field Service Representatives

The other side to continually testing and evaluating systems is the over reliance on field service representatives (FSRs) for many of these new systems. Some FSRs are barely more knowledgeable than soldiers on the equipment, but contracts required FSRs to perform installation, maintenance, and reloading operations. Soldiers, not civilian contractors, should be installing, fixing, and reloading equipment in combat—especially in large-scale combat operations. Yet reducing FSR reliance is easier said than done. We have all felt the pain of training a soldier to understand a new Army system, only to lose the capability to another unit before knowledge proliferates. Civilian FSRs are sometimes more permanent than military personnel and can be better resourced and trained to maintain and improve complex systems. Technical equipment requires formal education and hours of hands-on learning in the field, something most units struggle to sustain. We owe units a sound plan for fielding and maintaining C-UAS systems, and acquiring the skill to employ them effectively before they end up permanently parked in the motor pool.

The recent combat experience gained in Iraq and Syria remains the most significant US C-UAS fight to date even, if it will one day seem miniscule compared the future LSCO fight. Therefore, like a preseason matchup, any good sports team would use the experience to prepare for the long season ahead. And because the attacks 2/10 faced came repeatedly over a four-month period, across a variety of locations, each with different characteristics, the value of the empirical and quantitative data collected by the brigade makes the experience truly valuable to share with the force across operational domains.

As a reminder, the experience of 2/10 was purely a defensive fight. The challenges of performing C-UAS in the offense for a mobile fight are amplified. The current arsenal of missile interceptors, radars, directed-energy lasers, and nonkinetic defeat systems were only tested in fixed-site settings. The notion of establishing a layered air defense against Group 2 and 3 OWAUAS for major attacks like a wet gap crossing would require significant adaptations to the methods developed by 2/10 in OIR.

Still, static defense against OWAUAS will also remain relevant in future conflicts. Key theater support nodes like airports and seaports of debarkation will always serve as high-payoff targets within range of drone attacks. The UAS variants intercepted by 2/10 were often launched hundreds of kilometers from their intended targets. Attack drone ranges and payloads will continue to increase over time, so no matter how mobile frontline troops are in a LSCO fight, rear areas will need robust C-UAS defenses as well.

The C-UAS fight in CJTF OIR was a true example of what the chief of staff of the Army describes as transforming in contact. Consistent evaluation and testing by leaders across the tactical level was critical. Junior leaders found creative ways to isolate variables and test them. They could not just trust that the "green" status on a dashboard was good enough. They had to anticipate the next threat, then start adjusting and understanding their equipment. When select outstations began discovering new solutions, the unit organized weekly review sessions over secure teleconference video calls to spread lessons across the formation. It became a race for information as the enemy continually adapted against our defensive measures in real time.

The defense by 2/10 against more than 170 attacks from state-sponsored militia groups can help orient leaders to the challenges that the US military will face in future conflicts. Combat will continue to evolve, but the principles and insights outlined above are meant to accelerate the learning curve and highlight concerns for maneuver formations as our enemies push us toward a multidimension fight

Drone, Counterdrone, Counter-Counterdrone: Winning the Unmanned Platform Innovation Cycle

Zachary Kallenborn and Marcel Plichta

Shooting down drones is now an international pastime. In Ukraine, Russia, Sudan, Myanmar, and the Red Sea, militaries are scrambling to get their hands on counterdrone systems. In June, the US Navy issued a call for immediate kinetic counterdrone solutions and the UK is racing to have a high-energy laser operational as soon as possible. Market analysis estimates the global counterdrone market could reach $10.56 billion by 2030.

Global militaries, manufacturers, and pilots are not standing idly by.

Drone counter-countermeasures are a critical part of the competition between drone offense and defense. Today's drones can defeat countermeasures through a broad range of technologies and tactics. Drones might fly nap-of-the-earth to avoid detection, adopt greater autonomy to reduce the effects of jamming, fly in mass to overwhelm defenses, incorporate onboard defenses like antiradiation missiles, and more. Our new *Joint Force Quarterly* article, "Breaking the Drone Shield," describes and analyzes eleven such counter-countermeasures.

Drone warfare is best understood as a call-and-response innovation cycle, with each side responding to the other's innovation. Whether drone offense or defense dominates will vary based on advancements in technology, the degree to which adversaries adopt those technologies, as well as supporting doctrine, organization, training, leadership, personnel, and facilities. In the 2020 Nagorno-Karabakh War, drones gave Azerbaijan a significant military advantage because Armenia had limited defenses. Meanwhile, US forces in the Middle East are continually bombarded with low-cost drones from various terrorist and insurgent groups. Attackers are improving their drones too—Iran's new jet-powered Shahed-238, for instance, boasts increased speed, albeit with a higher cost—and doing so more and more quickly. This accelerating cycle of innovation may also manifest differently across domains as drones are increasingly used on land, at sea, and beneath the waves.

While much focus has been given to drones and drone countermeasures, little analysis has looked at the tactical and strategic implications of counter-countermeasures. US policymakers must correct this oversight to ensure US forces are equipped with the tools they need in an increasingly drone-saturated global operational environment. Counter-countermeasures should be based on a detailed understanding of adversary countermeasures because to break drone defenses, you first must know what defenses you're breaking.

Innovation and Counterinnovation

There are more ways than ever to defeat a drone—from radiofrequency and navigation system jamming to surface-to-air missiles, air defense guns, and plain old shotguns. RUSI estimated in 2023 that Ukraine was losing ten thousand drones per month, and this drone expenditure is likely matched on the Russian side of the ledger. For all the methods of downing drones, none is perfect. Drone countermeasures come with limitations and vulnerabilities that drone manufacturers can exploit. Even though the survivability of drones will define their utility in conflict, the subject receives scant attention from commentators.

Ukraine and Russia have used traditional air defenses to counter drones throughout the conflict. These have the advantage of already being fielded and understood by most militaries. Even supposedly obsolete air defense systems are finding a second life in a counterdrone role, such as the Gepard antiaircraft guns, whose production ended in 1980. The drawback is that many traditional air defense systems are not economical to use against small drones, particularly at range. Medium- and long-range air defense systems, like the Patriot, NASAMS, or S-400, are an order of magnitude more expensive to operate and resupply than all but the most expensive drones. Newer kinetic systems, like the L3 VAMPIRE, are less expensive but lack range, meaning militaries would have to procure and operate many of them to cover the same area as more advanced platforms. Attackers can turn these weaknesses to advantage. Russia's continual bombardment of Ukraine's infrastructure using mass drone attacks forces Ukrainian defenders to expend expensive magazines on low-value drones and keeps air defenses away from the front line to protect the rear. Even when Ukraine shoots down large percentages of incoming drones, civilians still feel the impact.

Naturally, both sides have sought more affordable, sustainable options like directed energy. Drones usually operate with either a wireless link to the operator and a link to a global navigation satellite system. Jamming or spoofing impedes the drone's mission. These systems have the advantage of being low-cost, requiring little more than the power necessary to

supply the jamming electromagnetic waves. Both sides in the war kicked off by Russia's invasion of Ukraine are slapping jammers on everything they can, including Russia's infamous "turtle tank."

The defensive emphasis on jamming resulted in an offensive emphasis on technologies and the techniques and tactics to counter them. Both Ukraine and Russia have been and are working to incorporate artificial intelligence, autonomy, and electronic defense. Basic countermeasures like frequency hopping on drone control systems have become commonplace. Both sides of the war in Ukraine began incorporating terminal guidance systems, in which drone pilots identify a target, and the drone matches the images or video it sees with the operator-selected target to make sure the drone is on target even if the connection is lost.
Increasingly, Ukraine and Russia are also focusing on navigation systems to accommodate temporary losses in satellite navigation signal, such as inertial navigation units and terrain mapping.

Of course, it is not necessary to wait until drones are in the air to counter them. Militaries are also pursuing offensive antiair operations. Large numbers of drones are stored or placed on airfields and often require numerous operators and support staff. Striking the drones or their operators on the ground reduces the pressure on counterdrone operations by eliminating threats before they emerge. In April Ukraine used a light aircraft converted into a drone to strike a facility where Russia was manufacturing a variant of the Shahed-136 for use against Ukraine's cities. A few months earlier in January, the United States struck twenty-eight locations in Yemen affiliated with the Houthi rebels' missile and drone program, including munitions depots, launching systems, and production facilities. Striking drones on the ground or the factories that produce them will prove to be a key part of counterdrone operations as militaries grow dependent on steady flows of drones to the front lines. Of course, this may create unexpected escalation risks if those facilities are located near population centers or other critical assets.

The Future Drone Fight

The back-and-forth between innovation and counterinnovation in drone warfare can be expected to evolve in new ways. Currently, focus is on two technological developments: directed-energy weapons and autonomy. Neither is immune to counterinnovation. New drone defenses like high-energy weapons will incentivize the creation of new measures to minimize their impact or target those systems. Autonomy will also force defenders to consider what systems to use and how to exploit autonomous systems for their own ends. Each will be discussed in turn.

The Department of Defense is quite interested in directed-energy weapons like high-energy lasers and high-powered microwaves as the future of drone defense. In April, the Army announced it had sent two such laser systems abroad to protect US troops. Lasers and microwaves promise to down drones at very low per-shot cost and reduce the risk of collateral damage by not intercepting the drones kinetically. Microwaves have the added advantage of a large area of effect that can knock down several drones at once, which will be useful as the scale of drone attacks increases and actors field drone swarms.

However, directed energy is no panacea. Lasers and microwaves come with trade-offs that create opportunities for adaptive tactics, techniques, and procedures. A significant weakness with both kinds of systems is that the effective range is generally short. Although the shorter range might be fine for point defense of strategic targets, the value may be limited for area defense. In addition, lasers typically require several seconds on target to create harm, and particulates in the air like rain or smoke can disrupt that. For example, American forces that tested a laser mounted on a Stryker found that the system struggled to function on a moving vehicle in tough conditions. To exploit these weaknesses, attackers might deploy drones during rainy or foggy weather, relying on the higher environmental hardiness of their drones. Although bad weather might also inhibit visual-based navigation and targeting systems on the drone, that may not matter much for static targets with known locations. The Office of Naval Research has several lines of research aimed at countering directed-energy weapons, including material hardening. Likewise, future antiradiation missiles may have seekers able to target directed-energy systems. That could be a big challenge. Although lasers and microwaves have low cost per shot, they often have high cost per system: the US Army contracted Epirus for four prototype microwaves for $66.1 million, or approximately $16.5 million per unit. If an adversary can affordably target and destroy those systems, the low-cost advantage may turn into a high-cost risk. Plus, directed energy often requires significant power, which may be disrupted or depleted. An attacker might use cheap decoys of plywood and foam plastic to burn through the system's stored power, before launching a larger attack.

On the other hand, autonomy presents its own set of challenges and opportunities. The Department of Defense is investing heavily in building autonomous drones, and Ukraine is reportedly working on machine-vision and terrain-mapping solutions to defeat Russia's extensive array of jammers. As autonomous drones become more ubiquitous and reliable, counterdrone methods like jamming might have less utility. Autonomy also opens new avenues for hackers or other actors to manipulate drones for nefarious ends. For instance, several years ago Chinese researchers tricked a Tesla

autopilot into steering the car into the oncoming traffic lane. The case raises the possibility that hackers could defeat or confuse autonomous systems into crashing, not recognizing a target, or even redirecting against its operators. As militaries adopt true drone swarms capable of autonomous communication the risk of interference will grow. The communication and collaboration necessary to create a true drone swarm may also create an opportunity for defenders to trick the whole swarm in such a way that errors propagate to every individual drone. More autonomous drones means more potential opportunities for failure as defenders learn to find and exploit the weakest link.

For defenders, autonomy might create dilemmas around incentives to shift to kinetic solutions, especially for homeland defense and security. Law enforcement and homeland security officials often rely on various jamming systems, including questionably effective handheld jammers. However, if drones rely on autonomous navigation, command, and control systems, then defenders will likely be forced to use physical effectors to shoot down or capture them. Although physical effectors vary in their risks from nets to surface-to-air missiles, the potential risk to bystanders can be expected to increase. If a major political leader were giving a speech and a potentially hostile drone showed up, would the Secret Service shoot it down over a crowd? Of course, the nuance is that even autonomous drones may differ in their immunity to jamming, as certain modern drones have automatic, difficult-to-disrupt return-to-home functions if the drone loses connection to its controller. But older, future, and do-it-yourself models may not have the same restrictions.

The Way Forward

Policymakers need to adapt to the back-and-forth evolution between drone offense and defense. In some contexts, drone offense might have the advantage while countermeasures are being developed, while in others effective countermeasures may blunt the impact of drones. That balance will vary in different locations, because actors will naturally differ in their access to various drone and counterdrone technologies and support capabilities. How domestic law enforcement protects airports and sporting events from drones will differ from how the military protects overseas bases.

To respond to—and get ahead of—the drone and counterdrone innovation cycle, American policymakers should consider three broad recommendations:

First, *understand the adversary*. Intelligence and defense officials study adversary drones intently, but they should give equal scrutiny to the US ability to down drones, and how quickly both sides of the equation are evolving. Close attention should also be given to nonstate actors, as insurgents, terrorists, and organized crime groups are also seeking to reduce the impact of drones on their operations. Simple jammers are often not difficult to make with off-the-shelf components, and are available commercially in some regions. Examining how adversaries conduct counterdrone operations should inform American drone development. That means investing in and expanding enabling capabilities like drone exploitation and forensics to quickly and effectively collect relevant information from downed adversary drones. Intelligence officials will need to work closely with acquisition officials to set appropriate requirements for industry. At the same time, balancing survivability with cost is key: hardening a drone against lasers may have little value against an adversary employing mostly kinetic defenses.

Second, *increase the pace at which the United States can develop and deploy counter-countermeasures*. The Ukrainian military claims a three-month innovation cycle to bring improvements to the battlefield in its war against Russia. That demonstrates that drone makers need to continuously develop solutions and modify drones to keep them survivable. The American development and procurement system is not optimized for such a rapid pace of evolution. The United States is in the middle of an agonizingly slow pivot from using drones in the permissive environments of the post-9/11 wars and counterterrorism operations to using them against adversaries with countermeasures. The slow pace has already cost the United States millions of dollars' worth of drones. For instance, the United States has reportedly lost multiple drones over Houthi-controlled territory in Yemen. The issue led officials to look at a few modular solutions, like the MQ-9's self-protection pod. Not every drone needs a full suite of expensive countermeasures, especially drones designed to be attritable, but drones will need to evolve faster than adversary countermeasures to consistently complete their missions.

As the United States and its allies ponder what would be needed for future confrontations with Russia, China, and Iran, focus should be on investment in research and development that can rapidly identify and target drone countermeasures, such as home-on-jam seekers and antiradiation missiles that operate on the microwave spectrum. In addition, the Department of Defense and each service should proactively investigate and experiment with tactics, techniques, and procedures. The White Sands Missile Range, Red Sands Integrated Experimentation Center, the UAE-based X-Range, and other test ranges already test and evaluate counterdrone systems. If they are not already doing so, the test ranges should incorporate and expand threats to counterdrone systems, and how American forces might attempt to defeat current and future adversary drone defenses. These results should inform virtual modeling and simulations to understand how changes to the probability of detecting or killing drones affect larger tactical and operational environments.

Third, *understand the trade-offs of counter-countermeasures*. Developing, manufacturing, and deploying counter-countermeasures could be expensive and raise unit costs. Hardening drones against every countermeasure can quickly lead to endlessly trying to keep pace with every new kind of drone defense. A main benefit of drones is affordable mass, so integrating expensive counter-countermeasures may diminish their core value to global militaries. Plus, counter-countermeasures may consume power, payload, compute, and other limited resources. So, the Department of Defense needs to carefully account for and consider these trade-offs in making acquisition decisions, organizing war games and exercises, and deploying and employing drones. In addition, the Department of Defense should support modular designs, and consider the appropriate drone fleet composition to account for regional and adversarial differences in drone defenses. Since different adversaries field different types of air defenses, acquisition officials should look for and incentivize modular solutions that can be added or removed for different kinds of missions. Effective solutions are likely to be technologically innovative with relatively low system complexity, so the Department of Defense would likely benefit from opening funding lines in the Small Business Innovation Research and Small Business Technology Transfer programs around specific counter-countermeasures of interest. Likewise, since American allies and partners are likely to face the same problem against similar adversaries, opportunities might exist for international development work, especially with states who already own American-made drones.

Drone innovation—like all military innovation—has been, is, and will be a continuous iteration between offensive and defensive innovation. The balance between US drones and adversary countermeasures will influence how future conflicts play out and US decision-makers must be sure US drones can complete their missions in nonpermissive environments. The United States must break the drone shield, and ensure it stays broken.

Small Units Need Protection from Drones—But What Capabilities Should a Light, Maneuverable Counter-UAS Platform Include?

Iain Herring and Gavin Berke

Imagine you are an infantry platoon leader, moving with your soldiers in a tactical formation toward your objective. Suddenly, indirect fire is raining down on your position. You have a plan to react to indirect fire, and you order your formation to execute the plan. Your soldiers are well trained and well led by their capable squad leaders, and they start to move, immediately and rapidly, from the impact area. But as you move, you realize the indirect fire is walking with you—your soldiers can't escape it. What you haven't realized is that there is a small unmanned aircraft system (UAS) observing your movement, allowing the indirect fire to follow you and your soldiers through the woods.

Now imagine the same scenario, except this time you have a mobile counter-UAS (C-UAS) system that can track and shoot UAS on the move. Once again, your platoon is engaged with indirect fire. And once again, your platoon has a plan and executes it on your order, Your light, maneuverable C-UAS vehicle can move with you, detect the UAS observing your platoon's movement, and neutralize it. Within a matter of seconds, the indirect fire ceases. Your platoon can safely regroup and continue mission.

Because of the C-UAS vehicle, traveling with and ready to support the platoon, the second scenario leads to mission accomplishment. Unfortunately, the first scenario is much more likely for Army small units today. C-UAS systems currently used are static or only semimobile, meaning they cannot move when C-UAS systems are operating. These systems proved sufficient during Operation Inherent Resolve—as we experienced while integrating airspace and countering UAS threats in support of the operation. But they will not fit with maneuver units' mission sets in large-scale combat operations. For that environment, the Army needs a truly maneuverable C-UAS platform for light maneuver units. This platform must be capable of detecting and kinetically or nonkinetically engaging UAS threats on the move and light enough to be air-assaulted or air-dropped from the back of fixed-wing aircraft. And it must be able to counter both sides of the future UAS threat coin: on one hand, increasingly affordable UAS will be fielded in growing numbers, potentially even as swarms, and must be engaged with electronic warfare (EW) effectors; on the other hand, there is the prospect of increasingly sophisticated UAS completely

resistant to EW that may therefore be best neutralized through kinetic interception.

With this in mind, what might an appropriate platform look like? Although components of an optimized system would need to be procured, many of them are already in service with the Army and could be adapted to meet future C-UAS needs while allowing for future innovation.

Vehicle Platform

Potential vehicles that could form the foundation of a light, maneuverable C-UAS system include:

- M1301 Infantry Squad Vehicle (payload 3,200 lbs., maximum gross weight 5,000 lbs)
- M1297 Ground Mobility Vehicle (payload 3,200 lbs., maximum gross weight 7,300 lbs)
- MRZR Alpha 4 Light Tactical Vehicle (payload 2,000 lbs., maximum gross weight 5,090 lbs)
- M3 RIPSAW Remote Controlled Vehicle–Light (tracked vehicle, predicted payload 3,000 lbs., maximum gross weight 16,000 lbs.)

Most of these vehicle platforms are just beginning to be fielded by maneuver units throughout the Army including the 101st Airborne Division (Air Assault). Depending on the final gross vehicle weight, some may be better suited than others for a C-UAS vehicle with air-assault or low-velocity air-drop capability. Ideally, this vehicle could accompany dismounted troops as close as possible to their objective area. Then it would remain close enough to provide protection while the dismounted soldiers conduct their tasks on the objective area.

Weapons and Effectors

C-UAS Air-Burst Weapon

A 25-millimeter or 40-millimeter grenade launcher with "smart" ammunition could be specifically tailored for C-UAS engagements. The projectile would detonate in flight, one to three meters before the UAS target. The range would be calculated by a laser range finder built into the scope or sight system. Upon pulling the trigger, a three-to-five-round burst of projectiles would fly toward where the UAS was predicted to be located, then explode at and around that location. The burst would create a cloud of shrapnel in case the UAS changes direction or speed. The shrapnel from the explosive would neutralize the UAS propulsion or flight control systems.

Electronic Warfare Effectors

An EW system is necessary to jam UAS signals. The Titan RF system with an amplifier, for example, may be well suited to disrupt or degrade small UAS. Any EW system included on the vehicle would not have to be co-located with or mounted on the platforms kinetic weapons. The proposed effective range for this system is greater than five kilometers.

Larger UAS Threat Engagement

An option to engage larger UAS kinetically is the FIM-92 Stinger Missile, which could be used to defend against any UAS threats below fifteen thousand feet with a large enough heat signature. Due to the Stinger's backblast area, the weapon would have to be mounted slightly above other systems on the vehicle and only operate within certain degrees of freedom relative to those other systems. This could reduce the need for a dedicated MANPADs (man-portable air defense system) operator. The effective range for this system is four kilometers.

Distance- or time-defined programmable detonation air-burst grenade rounds would be less expensive per munition and engagement than firing radar-guided missiles or hundreds of high-caliber munitions at UAS threats. Low-cost munitions and compact firing systems are an answer to system fatigue and munition supply constraints that currently limit C-UAS deployment in combat theaters.

A heavier solution for target acquisition and engagement could incorporate a system like the Common Remotely Operated Weapon Station (CROWS). A system like the CROWS could be a well-suited firing platform in conjunction with a remote-controlled vehicle platform like the RIPSAW. A system with these capabilities is still recommended to have man-in-the-loop engagement control to prevent fratricide of friendly or neutral UAS.

Detection

Before engaging UAS, either kinetically or nonkinetically, the C-UAS platform must first be able to detect UAS. This could be done passively or actively, but the detection capability will also be dependent on a reliable power source.

C-UAS Detection System Power

An UPS (uninterrupted power supply) battery pack that provides electricity for four hours or more with a tactical quiet vehicle-mounted generator should be used to maintain noise discipline. The generator would supply power to C-UAS systems while the vehicle is moving. The aim should be for the generator to provide eight hours per full tank of fuel to supply power for C-UAS systems and charging the UPS battery pack. The power output required of the generator is probably between three and ten kilowatts. The UPS battery pack would primarily be used while the generator is shut off. This design course of action increases unit survivability by reducing noise signature from combustion engines or generators near the objective area.

Passive-Active Combination Detection System

Passive radars could be used to detect UAS by measuring the change in electromagnetic frequencies made by the motion of a potential UAS. That detection must occur amid a wide range of commercial electromagnetic waves already present in an area of operation. Common sources and types of electromagnetic signals present in areas with infrastructure include Wi-Fi, cellular, civil radio and television, commercial satellite communications, satellite PNT (positioning, navigation, and timing), civil air radar, and weather radar. An excellent attribute of passive radars is that they do not reveal the observer's position because they do not emit radiation. This style of detection may be degraded in areas where commercial electromagnetic signals are not constantly present or in a scenario where sources of electromagnetic radiation do not have electricity.

Another sensor form of detection utilizes light detection and ranging, or lidar. It works by emitting and receiving laser pulses of nonvisible light that reflect off physical objects. The reflected light is received by the sensor and is converted to create a digital image of that object. Lidar could be used in conjunction with a passive radar to create high-confidence air tracks at short range. Lidar is commonly used on automobiles in the United States to detect other cars, pedestrians, or obstacles in a roadway. When used for this purpose, Lidar currently has an approximate range of 250 meters while the vehicle moves.

A system that fuses acoustic and optical sensors—multimodal unmanned aerial vehicle 3D trajectory exposure system (MUTES)—was tested at a distance of 480 meters by Siyi Ding, Xiao Guo, Ti Peng, Xiao Huang, and Xiaoping Hong, who published the results of their test in 2023. The conclusion of their report read:

Our results demonstrate that MUTES, which integrates a 64-channel microphone array, a camera, and a lidar, can provide wide-range detection (90° × 360°) and high-precision 3D tracking for UAVs [unmanned aerial vehicles]. A coarse-to-fine and passive-to-active localization strategy software was implemented in MUTES, with a well-designed microphone array capturing acoustic features and estimating the coarse position of the sound source, and the optical modules being used for further verification and tracking. Additionally, we trained an environmental denoising model to extract drone acoustic features, overcoming the drawbacks of traditional sound-source-localization approaches. A Kalman filtering algorithm for the fusion of three sensors proved to be effective and achieved the accuracy of RTK [real-time kinematic]. In terms of both hardware and algorithm, MUTES represents an innovative multimodal detection and tracking system.

This demonstrates a combination detection and tracking method of UAS can be effective. Future system combinations could use passive radar, Lidar, camera, or acoustic sensors.

Active Radar Option

An active radar capable of detecting group 1-3 UAS (the categories of the smallest systems) while moving is ideal to be prepared for future threats. The RPS-42 MHR, used with the M-SHORAD Stryker platform may be suitable for this vehicle. This would include multiple radars fixed at different points of the vehicle to provide real-time detection, early warning, and engagement capability for the maneuvering company or platoon. This could also allow for tracks to be pushed over a joint force tracker such as Link 16 to adjacent units and higher echelons. Optimally, the proposed detection range for this system is fifteen to twenty kilometers.

Positive Identification Enhancers

Multiple identification capabilities are highly recommended for the human-in-the-loop operator to identify UAS and prevent fratricide. Thermal and night vision are also necessary due to probable hostile threat windows and friendly maneuver operation timelines. The following systems have shown to improve soldier lethality with air defense systems: thermal sight, night vision, laser range detector for air-burst grenade programming, and a friendly UAS IFF (identification friend or foe) interrogator.

Network Connectivity

Tactical SIPRNet (Secure Internet Protocol Router Network) connectivity on the move would allow this system to send and receive air tracks, which would help filter out unknown air tracks. Furthermore, this would help with communication between multiple systems, making it easier to decide which system on the battlefield has the highest kill probability.

A mobile broadband kit or Starlink could be used to obtain connection to a tactical SIPRNet,. One Starlink system, one wireless router, one KG-175D TACLANE, two PacStar switches, and one four-port router could be configured to provide this capability, with all of the components fitting in a medium-sized rucksack. This network, in conjunction with a Forward Area Air Defense laptop, would allow tracks to be uplinked from the detection sensor on the light, maneuverable C-UAS vehicle to adjacent units and the higher command while also enabling tracks to be downlinked from higher echelons. This capability provides two functions: first, it provides commanders quick notice of air threats around their platoons or companies; second, it would provide even earlier warning to maneuver elements of threats in the airspace.

If these capabilities were packaged into a single vehicle light enough to be sling-loaded by a helicopter, detection could occur as far away as twenty kilometers, EW engagement at

more than five kilometers, and kinetic engagement within two kilometers. This would allow crews to see tracks approaching their positions and have enough time to warn adjacent units and their leadership, and then engage targets with EW and kinetic fires on the move. A lightweight solution like this could also provide advanced C-UAS capabilities to remote areas where a helicopter is the primary mode of entry.

To be sure, until such a system is built and fielded, the discussion remains entirely conceptual. But if the Army makes it a priority and works toward a light, maneuverable C-UAS vehicle that consumes minimal fuel and electricity, then it could be exactly what the platoon leader from the opening vignette needs to protect US soldiers from constantly evolving UAS threats and accomplish the mission.

The Return of Tactical Antiaircraft Artillery: Optimizing the Army Inventory for the Era of Small Drone Proliferation

Benjamin Phocas and Peter Mitchell

As the war in Ukraine has unfolded over the past two years, a seemingly endless series of videos have emerged from the battlefield, depicting drones—of various shapes and configurations—targeting and destroying personnel and equipment on both sides. Often, these systems are used at the microtactical level, to target individual soldiers and vehicles, and thus cannot realistically be defeated by current air defense systems. At the same time, across Iraq, Syria, Jordan, and the Red Sea, we are seeing the proliferation of one-way suicide drones being used against US, allied, and partner forces. To engage and destroy these drones, US forces are mostly reliant on two tools—shipborne weapons from a family known as Standard Missiles and Close-In Weapon Systems, which are essentially gatling guns attached to radars that throw up a wall of lead to defeat closing aerial threats.

The Phalanx Close-In Weapon System, with a maximum effective range of 1,500 meters and firing seventy-five rounds per second, has proved effective in countering modern drones. Its land-based variant, the well-known C-RAM, was similarly effective during the Afghanistan War against a variety of indirect fire. However, it is confined to a stationary defensive role, guarding warships and bases. On the modern battlefield, where drones are becoming prolific, there is a major gap that needs to be filled in counterdrone weapons systems.

The current US systems, and those in development, are effective under some conditions, but with the increased proliferation of cheap, lethal drones—and the threat of drone swarms looming on the horizon—having a system that can engage rapidly with lethal accuracy, in an extremely short period of time, is critical to defeating multiple fast-moving targets in close proximity.

The US Army should invest in modernizing its air defense artillery forces, to include dedicated antiaircraft artillery (AAA) batteries that are capable of defeating the threat posed by small unmanned aircraft systems (sUAS). Specifically, the service should reintroduce legacy-style gatling gun systems and work to field new systems rooted in the proven concept of the wall of lead.

The sUAS Threat

Many of the sUAS already present on today's battlefield are so small that they can be carried in one hand by their operators. In Ukraine, drone operators are working in coordinated hunter-killer teams, where one reconnaissance drone identifies targets for a team of attack drones to destroy. In a not-so-distant future, these hunter-killer teams could be further developed into advanced swarms that could operate semiautonomously, with a human only in the loop to approve targets prior to engagement, or even entirely autonomously of human operators. Many of the drones currently being weaponized are commercial, off-the-shelf systems designed for civilian use and, importantly, some are designed specifically for racing. These racing drones are extremely fast and nimble, often mounted with multiple rotors allowing them to rapidly move omnidirectionally. They are extremely hard to pinpoint and target, as they can move in an unpredictable, nearly insect-like manner in order to close with their prey.

After closing the distance, these deceptively small drones are capable of dealing devastating damage. Whether individual soldiers hiding out in craters and trenches or heavily armored T-series tanks, a single drone can find, fix, and finish targets with precision and efficiency.

Neither side of the Russo-Ukraine conflict has created a fully successful counter to this threat. Jamming systems have been effective, but often require a line of sight to cut or control the signals between a drone and its operator. In several documented instances however, these drones have been successfully downed by shotguns, machine-gun fire, and in at least one instance, a homemade gatling gun composed of a dozen AK-74s. The Gepard AAA system, given to the Ukrainians by Germany, has been highly effective in downing larger drones, such as the Iranian Shahed series that Russia has purchased and is now producing itself.

What this means is that on both sides of the war in Ukraine, the combatants have been forced to improvise and adapt under austere conditions—and have found success in simply throwing up walls of lead to down sUAS. This discovery has implications for the Army, the US armed service that will most likely face the brunt of the sUAS threat in a future large-scale combat environment. Put simply, it does not currently maintain a system that can effectively defeat these drones along the front lines.

M-SHORAD and Avenger Systems: Gaps in the Air Defense Inventory

The two current tactical air defense systems fielded by the US Army are the AN/TWQ-1 Avenger, which is mounted on a Humvee, and the new M-SHORAD (Maneuver–Short Range Air Defense) system mounted on the Stryker armored fighting vehicle. However, neither is suited for tactical and operational air defense in support of maneuver elements on the battlefield.

The first and most obvious issue is that both of these systems rely primarily on surface-to-air missiles to defeat targets. For targeting small, cheaply produced and converted sUAS, expensive missiles are simply not a cost-effective method of destruction. Additionally, these systems can only fire a small number of missiles (single digits for both platforms) before they are required to reload.

The M-SHORAD also mounts a single-barrel, 30-millimeter chain gun, similar to that on the Apache gunship. Even this weapon, though, is not suited to tracking and targeting small, fast-moving objects mounted with hand grenades, for example, or converted warheads for rocket-propelled grenades. It does not have the rate of fire to be able to throw up the mass of bullets necessary to defeat the aerial maneuvers of a drone, and certainly not if there are several of them. Footage in Ukraine shows that soldiers on both sides have attempted to use their rifles to defeat these drones, and they are rarely successful. It is simply too hard to hit such a small target with single accurate rounds.

Furthermore, these assets were not designed to be operated on or near front lines. Enemy drones operating in a swarm, with some dedicated to conducting suppression of enemy air defenses, could easily defeat an Avenger mounted on an unarmored Humvee or an M-SHORAD mounted on the lightly armored Stryker, only rated to stop 14.5-millimeter rounds. Neither could withstand direct hits, or potentially even near misses, from explosive-laden drones that have demonstrated the capability to destroy Russian tanks.

Additionally, as both these systems are wheeled, rather than tracked, they lack the same maneuverability and mobility of tracked vehicles, particularly in muddy terrain—like that found in Eastern Europe in the spring and in the Indo-Pacific region during rainy seasons.

It is important to note that these systems are both extremely valuable air defense assets that should continue to serve in rear-area aerial security roles. However, they do not offer the capability and protection required to counter the growing threat from sUAS. Nothing in the current US air defense arsenal has the protection necessary to operate near the front line and the fires capability to destroy swarms of cheap drones.

What's Old is New Again

The US Army was not always in such a predicament. In fact, until the mid-2000s, the service maintained an armored vehicle that could provide air defense along the forward line of troops. The M6 Linebacker was a modified Bradley Fighting Vehicle that simply replaced its turret-mounted TOW missile launcher with a launcher that carried Stinger missiles. The M6 also kept the Bradley's organic 25-millimeter chain gun for additional air and ground targeting capability. The M6 was fully capable of operating in a mechanized formation as an armored air guard that could maneuver and provide constant overhead protection simultaneously. However, similar to the M-SHORAD system, the M6 was also only equipped with a single-barrel cannon that fires too slowly to make it effective against small drones.

Thus, we must look further back into history to the predecessor of the M6, the M163 Vulcan Air Defense System. The M163 was an unsubtle bullet hose. It was little more than an M113 armored personnel carrier with a 20-millimeter Vulcan rotary cannon, similar to those mounted on the F-16 and A-10, inelegantly slapped onto the top. It was capable of firing a whopping three thousand rounds per minute in burst mode or one thousand rounds a minute cyclic mode, with rounds set to detonate at 1,800 meters.

The M163 was sold to the Israel Defense Forces, who modified the design, and created the improved Machbet variant, which added four Stinger missile launch tubes to the Vulcan cannon for targeting a variety of threats.

The M163 had its major drawbacks too. It lacked an organic radar system and relied on human gunnery to acquire and target enemy air assets. The M113 vehicle it was based on is also limited, primarily in the fact that it is a more lightly armored personnel carrier, not designed to withstand the same level of fire as tanks or the more modern Bradley. A new system, the M247 Sergeant York, was planned for development in the 1970s and early 1980s, but the program was an utter debacle and was scrapped in 1985.

Both of these systems once in the US Army inventory, the M6 and the M163, offered something missing today. They both had the advantage of being tracked vehicles, for instance. But each also had its deficiencies. The M6 had the armor but not the right firepower, while the M163 lacked the armor but packed the right punch, particularly in later variants. If the strengths of these two systems could be married, however, there could be an air defense vehicle with both the armor and firepower to operate alongside maneuver formations and able to defeat both sUAS and larger threats such as helicopters.

The Solution

The US Army must invest in a mobile air defense system with the capability to effectively defeat the enemy sUAS threat, while also retaining the protection and maneuverability to operate in frontline areas.

The solution does not need to be a revolutionary system. Nor should it be. The threat from sUAS is here now, and a project that spends the next decade in research and development will not counter the present threat. The relatively simple and much more low-cost solution is to use older-model Bradley Fighting Vehicles no longer in active US service—there are nearly three thousand currently sitting in storage—and convert them to basic but functional AAA systems. These conversions would not require the invention of an entirely new vehicle platform and would only require an off-the-shelf existing system such as the Close-In Weapon System or the development of a similar, but more tailored, AAA system. Having such a system mounted on a vehicle that can operate under the dangerous conditions of frontline combat and is able to withstand all but a direct hit from an antiarmor-equipped drone or other weapons system, could be the difference between life and death for US soldiers in a conflict in the very near future.

Without such a system available, US ground forces will be vulnerable to attack from sUAS and will have no effective defense other than firing wildly into the air as an untold number of now deceased combatants in Ukraine did, to no avail.

Disrupting Systems: Cyber and Electronic Warfare

From Georgia to Ukraine: Seventeen Years of Russian Cyber Capabilities at War

Ketevan Chincharadze

In August 2008, as Russian tanks rolled into Georgia's Tskhinvali Region, not self-proclaimed South Ossetia, Georgian government websites were under cyber siege. Distributed denial-of-service (DDoS) attacks, defaced portals, and data theft disrupted communications as Georgian officials tried to urgently reach Western leaders, some on vacation, others attending the Beijing Olympics opening ceremony.

For the first time in history, a state had unleashed coordinated cyberattacks along with military operations. In post-Soviet, developing Georgia, with limited digital infrastructure and nascent social media, the attacks received little public attention and had minimal impact on combat operations. Seventeen years later, however, technological advancement and growing digital dependency have dramatically amplified the scale of cyber threats. The ongoing war in Ukraine illustrates this trend.

Russia's Cyber Experiment in Georgia

In the weeks leading up to the Russo-Georgian War, Russian hackers attacked Georgia's digital ecosystem to sow chaos within the Georgian government and society as Russian troops were amassing along the northern border. This marked the dawn of modern hybrid or gray zone warfare, which blends conventional military force with unconventional tactics, such as cyberattacks.

In July 2008, millions of DDoS requests overwhelmed Georgian websites in an attempt to disable both government and civilian servers. Close to the invasion, hackers began using techniques such as SQL injections, a more advanced assault, which enables attackers to bypass website protections and directly penetrate servers with malicious queries.

Numerous websites were defaced, and some even used photo manipulations to compare Georgia's then president Mikheil Saakashvili to Adolf Hitler. Hackers targeted key political, governmental, and financial platforms, including the websites of the Georgian president, the National Bank of Georgia, and the Ministry of Foreign Affairs. They also exploited lists of public email addresses and infiltrated government networks to extract potentially sensitive information.

Experts have suggested that Georgian internet traffic was rerouted through Russian telecommunications firms, whose servers also hosted malware used in the attacks. Additional evidence indicates that attackers manipulated an informal online poll on CNN's website to portray Russia's combat operations in Georgia as a legitimate peacekeeping mission. Russian bloggers then rapidly spread the poll across the country, urging their readers to visit CNN's website and select the response supporting Russian intervention. As a result, 92 percent of predominantly Russian participants voted in favor of the peacekeeping narrative before CNN ultimately removed the poll.

In 2008, according to the World Bank, only 10 percent of the Georgian population used the internet, compared to 82 percent in 2023. With such limited public reach at the time, the attacks were primarily aimed at demoralizing the government, diverting attention from military operations, and stealing intelligence. However, as internet access expanded across the country, so did Russia's influence on the public.

Moscow started using disguised, low-profile content to subtly shape public opinion and obtain user data without informed consent. In 2020, for example, Facebook removed News-Front Georgia, a Kremlin-linked outlet that had been actively spreading pro-Russian and anti-Western sentiments through an organized network of inauthentic accounts. According to the International Society for Fair Elections and Democracy (ISFED), the network included twelve fake profiles that disseminated pro-Russian content in thirty-one Facebook groups with over 521,000 members, in a country of just 3.7 million.

ISFED also uncovered twenty-six fake Facebook accounts and pages, disseminating Kremlin-backed Sputnik Georgia's content across forty-one public groups, reaching 1.2 million users. The operation used so-called *soft* content, such as posts about gardening, astrology, or local celebrities, to build trust with users before inserting links to Sputnik's articles with Kremlin-aligned narratives.

Russia has been steadily expanding its overt and covert operations since 2008. The annexation of Crimea and the ongoing war in Ukraine further demonstrate Moscow's continued advancement of its digital arsenal for modern warfare.

Advanced Cyber Operations in Ukraine

Strategically, the Kremlin began small in Georgia and significantly scaled its military and cyber warfare in Ukraine. The 2008 rudimentary attacks were an experimental foundation, evolving into broader assaults on Ukraine's communications and energy sectors in 2014, and ultimately escalating into a global threat targeting Ukraine and its allies during the full-scale invasion. However, much of Russia's strategy still follows a familiar playbook first tested in Georgia.

Just like in Georgia, Russia's first wave of cyber operations predated the 2014 annexation of Crimea. The attacks on the information systems of Ukrainian state institutions and private enterprises came during the 2013 mass protests that would become known as the Maidan Revolution. In mid-2013, Operation Armageddon targeted Ukrainian government, law enforcement, and military officials to steal sensitive information through phishing emails that tricked victims into clicking malicious links. Just three days before the Crimean status referendum, on March 13, 2014, Russia launched an eight-minute DDoS cyberattack on Ukrainian computer networks and communications to distract public attention from its military presence in Crimea.

Unlike in the Russo-Georgian War, Russian cyberattacks extended beyond the annexation of Crimea. In 2015, Ukraine experienced two assaults on three regional power distribution entities, also known as oblenergos, which impacted approximately 225,000 customers. In a US context, this would be proportionate to attacking the Omaha Public Power District, the Nebraska Public Power District, and MidAmerican at the same time. The US Cybersecurity and Infrastructure Security Agency concluded that the oblenergo "unscheduled power outages" were perpetrated by "Russian nation-state cyber actors."

By 2015, researchers had identified two prominent Russian hacking groups involved in Russia's cyberattacks against Ukraine: APT29 (also known as Cozy Bear, Cozy Duke, or Nobelium) and APT28 (also known as the Sofacy Group, Tsar Team, Pawn Storm, or Fancy Bear). These groups also played an important role during Ukraine's full-scale invasion in 2022.

Following the pattern established in 2008 and 2014, Russian hackers intensified reconnaissance efforts in Ukraine far ahead of the invasion. This included actions by APT29, which has been linked to the SVR, Russia's foreign intelligence service. In the lead-up to the invasion, government and university websites were defaced, spear-phishing campaigns targeted the energy sector, and DDoS attacks hit the Ministry of Defense and major banks. At the same time, coordinated disinformation campaigns portrayed Ukraine as an oppressor of the Russian-speaking majority in the country's east, echoing the CNN poll manipulation in 2008 aimed at framing Russian troops as peacekeepers in Georgia's breakaway South Ossetia.

Hours before the invasion, GRU Unit 74455, also known as Sandworm, the same Russian military intelligence group behind the 2017 NotPetya attacks, deployed a wiper malware called FoxBlade against Ukraine's digital infrastructure. Victor Zhora, a prominent Ukrainian cybersecurity official, called the attack "a really huge loss in communications in the very beginning of the war."

Hacking communication infrastructure to gain a military advantage is central to Russia's war strategy in Ukraine. In the weeks following the Sandworm incident, Russia made another attempt to shut down Ukraine's internet access by targeting three major telecommunications providers—Triolan (March 9), Vinasterisk (March 13), and Ukrtelecom (March 28). SpaceX's early delivery of Starlink terminals helped restore communications across Ukraine, and Russian forces quickly responded by trying to hack, jam, and disrupt Starlink's operations, though with limited success.

Russia's assaults also extended beyond Ukraine's borders to its allies—a significant step up from earlier practices. According to Microsoft, "By mid-2021, Russian actors were targeting supply chain vendors in Ukraine and abroad to secure further access not only to systems in Ukraine but also NATO member states." This practice intensified as the war escalated.

A coordinated cyberattack on Viasat satellite modems disrupted satellite communications across Ukraine and parts of Europe on February 24, 2022, the day of the invasion. The operation crippled Ukrainian communications, including internet access for thousands in Ukraine, and disrupted the KA-SAT satellite internet service across Germany, France, Hungary, Greece, Italy, and Poland. In Germany alone, more than 5,800 wind turbines were affected due to the loss of satellite connectivity. The EU publicly linked the attack to Russia.

Microsoft reported an increase in Russian cyber espionage throughout 2023 in at least seventeen European countries. It

also identified a new GRU-linked threat actor, Cadet Blizzard, active since February 2023, which targets organizations in Latin America and Europe, particularly in NATO countries supplying military aid to Ukraine.

From Georgia in 2008 to Ukraine in 2022, Russia transformed its cyber experiments into a sophisticated global threat. In July 2022, eight distinct Russian malware strains were deployed to breach forty-eight Ukrainian government agencies and enterprises, averaging two to three attacks per week.

Since the war, Moscow has used nine new families of wiper malware and two new ransomware variants, targeting more than one hundred Ukrainian government and private sector entities, including the Prestige ransomware, deployed in October 2022 in Ukraine and Poland. By late April 2022, Microsoft had recorded 237 cyber operations targeting Ukraine, including destructive attacks, service disruptions, espionage efforts, and coordinated disinformation campaigns.

But this is just the tip of the iceberg. In 2023, Shane Huntley, a senior director of Google's Threat Analysis Group called Russian cyber operations "aggressive" and "multi-pronged," while the general manager of Microsoft's Threat Analysis Center, Clint Watts, cautioned that Russia was continuously innovating with new malware. Further reports indicate that the Kremlin complements these cyber operations with extensive disinformation campaigns, blaming the West for the war in Ukraine and pushing pro-Kremlin narratives through more than one hundred thousand social media pages and Telegram channels.

The ongoing war in Ukraine, so far, represents the most vivid example of how cyber capabilities can complement activities in other warfighting domains. However, the overall impact of cyberattacks on Russia's ongoing war in Ukraine is still uncertain.

Can Cyberattacks Win Wars?

Despite the significant expansion of Russia's cyber operations from Georgia to Ukraine—even earning Russia a reputation for having some of the world's most formidable hackers—cyberattacks have not yet had a decisive impact on the war in Ukraine.

When Russian forces wanted to disrupt civilian infrastructure, they routinely bombed hydroelectric plants and other critical energy and water facilities across the country. In March 2024, Russia launched eighty-eight missiles and sixty-three Iranian-made Shahed drones against Ukraine's largest dam, leaving over one million people without electricity. The most severe internet disruptions have also resulted from such missile strikes rather than cyberattacks on Viasat satellite modems. Therefore, conventional kinetic operations continue to dominate Russia's operational approach to warfighting.

The cases of Georgia and Ukraine, however, show that cyberattacks can effectively disrupt government operations and sow uncertainty even if they do not yield decisive results on the battlefield. Disinformation campaigns can further sway public opinion in Russia's favor. The Kremlin effectively exploits the lingering anti-Western sentiments of Cold War generations, who still represent a significant portion of the electorate and political leadership in the post-Soviet zone. These sentiments are also reflected in a recent Friedrich Ebert Stiftung survey, which shows that more than a quarter of Ukrainians blame the United States for the war, 15 percent blame the EU, and 66 percent think Ukraine should avoid international involvement.

Russia, of course, is not alone in developing advanced cyber capabilities. Other major powers are closely observing and learning. China, for instance, is augmenting its government-based cyber arsenal and hiring a private network of hackers, which can be a growing threat to the United States. The US Justice Department has already charged twelve Chinese contract hackers and law enforcement officers for their involvement in global computer intrusion campaigns in March 2025.

In less than two decades, wartime cyber operations have evolved from rudimentary disruptions to sophisticated attacks on critical infrastructure and coordinated efforts aimed at undermining Ukraine's defense capabilities. While the digital domain may not yet determine the outcome of war, it has increasingly blurred the line between civilian and military targets—from disinformation campaigns targeting ordinary citizens to espionage infiltrating government institutions. And because cyber operations do not begin or end when the shooting does, this is the war front that never really goes offline.

Commander's Intent for Machines: Reimagining Unmanned Systems Control in Communications-Degraded Environments

Matthew Corbett

The Russian invasion of Ukraine has dramatically shifted the way the US Army envisions the employment of unmanned systems. Unmanned aircraft systems—drones—reportedly account for two-thirds of Russia's daily losses, surpassing all other weapon systems. First-person-view drones, small aircraft that require a human operator to essentially drive an explosive-laden system into a target, have become popular as cheap and effective weapons. As a direct result, countries around the world are scrambling to develop systems to both defeat them and adapt their arsenals to integrate them.

To do so, jamming has been employed against the small systems most prevalent on the battlefield. This involves disrupting the electronic control link between the drone and its human operator. Severing this link denies the direct connection needed for the operator to guide the system along its terminal path to an intended target, rendering it ineffective.

The prominence of jamming on the battlefield exposes a fundamental problem—the vulnerability of the link between a drone and its human operator. In all but a few systems, human operators are required for flight control, terminal guidance, payload and camera operations, and navigation. To mitigate this, autonomous systems are fielded in smaller numbers. However, these systems still require human input in the form of a rigid mission plan, and do not have methods to modify their tasks in the face of a changing mission without explicit human input.

This is the challenge, then, for military forces like the US Army that seek to leverage the capabilities of unmanned platforms while reducing their vulnerability to jamming: creating a method to provide intent to autonomous systems. The Army already has a model for this—the doctrinal concept of commander's intent. This is the method by which a commander gives subordinates a required end state and key tasks, and it becomes especially important when original mission parameters become untenable or communication with a higher headquarters is lost. Providing these machines with a version of commander's intent, coupled with AI systems to parse and create machine-readable tasks, offers a method to overcome vulnerabilities to jamming with minimal human oversight. While this idea does not necessarily remove the human operator from lethal kill chains, it does allow for more flexible and redundant drone employment in the hostile environments of the future.

A New Type of War

At the beginning of the war in Ukraine in 2022, drones were a part of ground combat in limited numbers. By 2024, both Ukraine and Russia were fielding drones capable of a myriad of tasks, both lethal and nonlethal. The production of hundreds of thousands of these systems has become a top priority for both countries, with Ukraine setting up domestic mobile production shops that specialize in the efficient manufacture of the smallest versions of these systems. As traditional weapons such as manned aircraft and armored vehicles have been attrited or have become too vulnerable to be tenable, these systems have replaced them as the most lethal weapon of this conflict.

As each side responded by jamming adversary drones, an ever-evolving game of cat and mouse emerged, leading to the introduction of so-called tethered drones. These systems create a physical connection between the system and its operator and reduce the possibility of electronic attack on the control link. As evidence of the use of these systems, recent pictures have emerged of treetops in the contested areas of Ukraine littered with fiber-optic cables from active or previous drone use. More recently, Ukraine has reportedly fielded semiautonomous drones capable of lethal strike to defeat Russian jamming.

Commander's Intent: A Model for Mitigating Control Link and Human Vulnerabilities

Both the Ukrainian Armed Forces and Western militaries have identified that rigid orders and inflexible tasks create unnecessary casualties on modern battlefields. In lieu of a direct connection to a higher command, soldiers must be able to create new tasks at lower echelons to achieve the overarching original goal of the higher headquarters. In the US Army and others worldwide, the overall intention for the end of an operation is communicated as the *commander's intent*.

In the US Army, commander's intent "provides a unifying idea that allows decentralized execution within an overarching framework." To repurpose this concept for unmanned systems, one subcomponent of commander's intent is the most important: key tasks. Key tasks are "significant activities the

force must perform" to achieve a desired end state, the terminal condition of a given mission. Key tasks can include effects on the enemy force or required friendly conditions. These discrete, well-defined tasks are the most easily translated into the tangible machine inputs we desire. Other elements of intent are also relevant. Some types of the commander's critical information requirements, for example—like the location of particular enemy systems or the location of friendly assets—can be preprogrammed before launch.

Aerial Reconnaissance: A Case Study

Key tasks can take many forms, some of which are not easily translated into discrete tasks for a machine. However, as a case study, we examine aerial reconnaissance, a common drone mission set, to explore the possibilities of using this information. Consider a fictional scenario in which a ground force is tasked to search for and destroy key enemy systems, identified in earlier planning and listed in order of importance (e.g., a high-payoff target list). These systems have been assumed to be around the unit's battlespace, identified in specific areas called target areas of interest (TAIs). The unit has been assigned the tasks of finding these systems and producing fire missions for subordinate artillery or other fires platforms.

The ground unit intends to use its assigned small, aerial drone platforms, each equipped with cameras and potentially lethal payloads. The unit assumes a position near the expected location of the targets and launches the drones. However, due to the impact of jamming and the lack of a physical tether, these systems are quickly isolated from their human operators. The operators can no longer give flight control commands or operate the drone's attached camera or weapon systems. Luckily, the drones are equipped with a method to respond to such a condition. They can adhere to the key tasks for the mission because they have been preprogrammed to revert to these tasks if their human counterparts are no longer available.

The drones move to their preprogrammed TAIs in sequential order. Each is equipped with modern aided target acquisition systems, using AI techniques like advanced object detection to scan incoming video frames for a particular set of objects, like the enemy systems identified on the high-payoff target list. If the enemy has also degraded their ability to locate themselves using GPS, as is often the case in Ukraine, these systems can do so using onboard inertial movement sensors and reverse mapping of their environment (e.g., simultaneous localization and mapping, or SLAM, a common technique for localization of mixed reality devices). This is roughly similar to the way the US military teaches land navigation.

If, at these TAIs, the drone's aided target acquisition system identifies an object or position that is classified by the same name as one or more of the items on its high-payoff target list, the system can produce a military grid coordinate based on its estimated location (assuming a GPS-denied environment) and a rough trigonometric transformation. All of this, however, is useless if the target information cannot be translated into a fire mission. The drone now requires a connection to its human counterpart to complete its assigned mission. Using the same inertial navigation from before, the drone can move toward the initial area where jamming severed its control link, and along the vector toward its operator, until such a link is reestablished and the target location can be transmitted. The drone is then retasked with the same intent or a new version as needed.

Technical Challenges

The underlying technical capabilities needed to adopt commander's intent for drones exist, but are currently unavailable in a single military unmanned system. Ukraine has fielded a lethal, autonomous version already. While its full capability is unknown, open-source information about it shows that a weaponized, drone-borne object-detection capability is already a reality. Encoding written intent as machine-readable tasks is possible as well. This is necessary to extract the key tasks and end state that may be especially useful in the commander's intent. Existing research already has shown that natural language processing can extract the semantics of written intent and translate this into its component elements and tasks. Beyond this basic and general research, little work exists that is focused on a military context. Finally, localization and navigation from known points without the aid of GPS satellites is an active field of research, but no research exists to leverage this technique for object detection and aided target acquisition. This is required for autonomous navigation in the GPS-denied environments that exist today and will likely exist in the future. Synchronizing this existing work with an eye toward military use cases is still necessary to fully implement intent in these machines.

Ethical Implications

Our fictitious drone from the story above was not required to use its lethal payload to destroy the target it located. Had such a need arisen, existing policies would govern the use of force by autonomous systems (or, in our case, a semiautonomous system operating in an autonomous mode). In the United States, current policy dictates that these systems are designed to allow their operators to "exercise appropriate levels of human judgment over the use of force." Using lethal force without explicit human approval is still a matter of debate and concern worldwide, and it raises questions about the uncertainty of identification, especially in civilian-populated areas. However, nonlethal drones, or lethal drones undertaking

an exclusively nonlethal mission, are less of a problem. In our example, unless that drone's key task was to *destroy* the identified target on the high-payoff target list, direct and lethal force was never required. Indirectly, however, the system relayed coordinates, likely resulting in lethal action taken by other units.

While ethical discussions surrounding this selective autonomy are likely to involve slippery slope arguments, countries like Ukraine and Russia have already begun developing and fielding autonomous systems. If more countries decide to follow suit, the question of how to guide and safeguard autonomy will remain. Specific, close-ended tasks, like those already given as part of the commander's intent, are an excellent option to mitigate these concerns and keep humans a part of the kill chain.

While technically feasible, current unmanned military systems are generally not designed to operate without direct human connection in dynamic environments. As more unmanned platforms with increasingly autonomous capabilities are developed, the challenge shifts from a technological one to one of employment techniques. Commander's intent, a well-known and tested method of issuing guidance in uncertain and communications-restricted situations, is a promising avenue on which to pursue a solution to this employment challenge. By predeploying key tasks and end state in the form of discrete instructions, as well as the augmentation of existing drones with autonomous navigation and aided target acquisition systems, US Army units will be best positioned to leverage drones' unique capabilities to accomplish missions in the face of future jamming environments.

Advantage Defense: Artificial Intelligence at the Tactical Cyber Edge

Zachary Szewczyk

In 2019, Rudy Guyonneau and Arnaud Le Dez captured a common fear in a *Cyber Defense Review* article titled "Artificial Intelligence in Digital Warfare." "The question of AI now tends to manifest under the guise of a mythicized omniscience and therefore, of a mythicized omnipotence," they wrote. "This can lead to paralysis of people fearful of having to fight against some super-enemy endowed with such an intelligence that it would leave us bereft of solutions." With the release of ChatGPT in 2022, it looked like that fear had come true. And yet the reality is that AI's use as an offensive tool has evolved incrementally and not yet created this super-enemy. Much of AI's real value today lies in the defense.

As Microsoft and OpenAI recently explained, today we see threat actors using AI in interesting but not invincible ways. They found five hacker groups from four countries using AI. At first, the groups used large language models for research, translation, building tools, and writing phishing emails. Later, Microsoft saw the tools suggesting actions after a system had been hacked. Although some argue that modern models could take on more, that seems premature. In stark contrast to fear that AI would unleash a wave of robot hackers on the world, these actors used it for mundane tasks. Defensive cyber forces, on the other hand, could use AI technology that exists today to meaningfully improve cyber defenses in four key ways: accelerating the pace of analysis, improving warning intelligence, developing training programs more efficiently, and delivering more realistic training scenarios.

First, endpoints and network sensors create billions of events per day across the Department of Defense Information Network. Today, "data overload" is not just a theoretical danger. It is a given. As Guyonneau and Le Dez pointed out, though, volume is only half the battle. Cyber analysts must also grapple with "techniques and strategies [that] evolve at a frantic pace, the former through the exigence imposed by early experiences in the field and the rate of technological development, the latter as our understanding of the stakes grows." It is not just the volume of data in the fifth domain that confounds understanding, but its complexity as well. This ocean of uncertainty is a prime target for two of the most common forms of AI, machine learning and large language models.

Machine learning won't turn data into knowledge by itself, but it can speed up analysis. These models might not know *why* an endpoint acts the way it does, but they can spot weird activity. At scale, they shift the burden of sifting through millions of logs onto a computer. As a result, people spend less time searching for the digital needle in the cyber haystack and more time on complex investigations. The challenge of training, tuning, assessing, using, and parsing the output of these algorithms, though, means that few use them well, if at all. Large language models can help. ChatGPT or the open-source Llama 3, for instance, can handle these tricky steps. Instead of coding a support vector machine, I can ask ChatGPT to "Build a support vector machine with this sample data." Instead of

sifting through pages of documentation to tune hyperparameters, I can ask Llama 3 to tune them. Tasks that once took data scientists hours can now take an eager analyst just minutes.

Large language models could also accelerate the pace of analysis as the backbone for analyst support tools. Cyber analysts start many investigations based on opaque alarms. For example, an alert that "Trojan:Win32" malware could have infected an endpoint might entail hours of work just to gather basic information. A large language model could instead create a brief report that explained the alert, assessed suspect files, collected facts about the host that raised the alarm, and offered next steps for the investigation. The prominent threat hunting and incident response firm Red Canary already does this with what it calls "GenAI agents." Externalizing mundane tasks like these would drastically accelerate the pace of analysis.

As a stepping stone between manual and semiautonomous investigations, one of my projects used large language models to build analyst playbooks. These playbooks guide junior analysts to approach complex investigations in a similar way as their more experienced counterparts do. They promote analytic rigor. The process of researching, understanding, and then creating detections and investigation strategies for such a vast array of malicious activities, though, takes months. Over the years I have seen many pursue this lofty goal yet inevitably fail. Using large language models and a bit of Python, though, I built a library of over six hundred playbooks—one for each technique in MITRE's ATT&CK matrix, a taxonomy of malicious actions in the cyber domain—in a few hours.

Second, machine learning could also help derive meaning from internet-wide scanning data for improved warning intelligence. The intelligence cycle has struggled to keep pace with the cyber domain. Many reports on servers used to launch attacks or control malware implants, for example, arrive far too late to do any good. They provide interesting but seldom actionable information. By finding the traits of those servers from internet-wide scans and training machine learning models to spot them, cyber analysts can use these tools on live data feeds to quickly find new malicious servers. Rather than acting on similar insights in days or weeks as reports make their way out of an intelligence cycle, this approach would operationalize intelligence at machine speed.

Third, AI could better prepare analysts for defensive cyber missions. Training, for instance, takes a lot of time and is hard to do well. I dealt with this just last year in the new 3rd Multi-Domain Task Force. Assigned to an Army service component command rather than part of the Cyber Mission Force, the unit's large cyber formation stood up without access to the training to do its job or any plan to obtain it. We found ourselves again re-creating the wheel by building our own training program. We planned to spend over a year on this project. After some experimentation, though, we found a way to use large language models to create the entire curriculum—to include lessons plans, training material, and even some hands-on exercises and assessments—in just a few hours.

Finally, AI could also improve hands-on training. Realistic scenarios are exceedingly difficult to build, run, and maintain. So much so, in fact, that they do not exist. Michael Schwille, Scott Fisher and Eli Albright recently described the challenges they faced when they tried to implement data-driven operations—using real-world data—into an Army exercise. As Guyonneau and Le Dez pointed out in their 2019 article, though, "If the corresponding data exists and can be acquired, a cyberteammate has the capacity to simulate any type of environment, whether friendly, neutral, or adversarial." An AI agent can handle almost everything. Where an entire team would have manually setup cyber ranges, an agent could generate code describing that cyber range and then deploy it in a common industry practice called *infrastructure as code*. An agent could also run realistic scenarios with synthetic actors that respond to trainees' actions in real time. No longer must analysts suffer through small, contrived events based on canned scripts put on by under-resourced training cells.

There are valuable roles for AI to play in cyber operations. As Jenny Jun recently described it—with admirable brevity—the effects of AI in the cyber domain will be "sharper swords, tougher shields." On the offensive side, though, those roles remain small for now, as Microsoft and OpenAI observed, and ultimately might not make offensive cyber operations relevant at the tactical level. Much of AI's value today lies in defensive cyber operations. As a cyber analyst, I have access to hundreds of billions of new records per day—a prime target for machine learning. When paired with improved warning intelligence, also through machine learning, this technology presents an opportunity to drastically reduce the amount of time threat actors go undiscovered—or even neutralize a campaign before it starts. An analyst support tool, built on top of a large language model, could further accelerate my pace of analysis. In the lead up to those operations, AI could help lessen the crushing burden of building and running training. Unlike many lofty ideas that over-promise and under-deliver, these goals are realistic and achievable with the resources line units have today. We say we want innovation. Here is the opportunity; we must seize it. This is how we can move toward meaningful use of AI at the tactical cyber edge.