SPECIAL COMPETITIVE STUDIES PROJECT

Defense Paper Series

Reimagining Command and Control with Human-Machine Teams

Lieutenant General (ret.) Clint Hinote, U.S. Air Force

This paper series includes discussion papers written by SCSP advisors to accompany our Defense Panel's 2024 Working Group Meetings. The views and opinions expressed in this newsletter are solely those of the authors and do not necessarily reflect the views or positions of SCSP.

The challenge of commanding and controlling military forces today is ageless, informed by centuries of historical experience and simultaneously unprecedented in its contemporary complexity.

From one perspective, the challenge of commanding and controlling military forces extends back to the beginning of organized warfare. As military historian Martin van Creveld writes in *Command in War*, "The problem of commanding and controlling armed forces, and of instituting effective communications with and within them, is as old as war itself." This historical backdrop offers invaluable lessons and doctrines. Fundamental elements such as Clausewitz's concept of *fog* underscore the need for clarity and decisive actions within the command structure. Effective prior planning allows commanders to prepare for various contingencies, preparing the commander and staff for changes through an appreciation of the overall military problem. Furthermore, logistical planning and execution make military campaigns possible, ensuring that forces are adequately supplied and able to act in concert. The principle of simplicity emphasizes the need for clear and concise communication in command directives, ensuring that strategies are easily understandable and executable by all ranks. As these were important considerations in the past—sometimes making the difference between victory and defeat—so will they be in the future.

The act of command and control (C2), a fundamental concept in military operations, refers to the exercise of authority and direction by a commander over assigned forces in accomplishing a mission. It is deeply interwoven with the very identity of the military commander. When a commander accepts the responsibility of leading a unit, they not only take on a title but also the implicit expectation that they will effectively guide the unit's actions, demonstrating both authority and insight in the process. In Western military culture, considerable emphasis is placed on comprehensive education and rigorous training to equip commanders with the skills necessary for effective command and control. This extensive preparation includes sizable investments in personnel and advanced technology, ensuring commanders are well-versed in leadership principles and technological tools to communicate directives to subordinates effectively. When enabled by these foundational elements, commanders can draw upon historical experience and military doctrine to execute successful operations.

Despite the historical context, the present era poses novel challenges in command and control that lack historical parallels. One significant factor is the evolving nature of warfare, where commanders must cope with an increase in scale, scope, and speed. A potential conflict between major powers, such as the United States and China, may not only be global in its dimensions but may also transcend traditional battlefield boundaries, affecting various sectors of society—including economic and social arenas. The weapons deployed in such a conflict would have rapid and widespread effects, creating a complex interplay of capabilities that demands an adaptive command structure. While history has witnessed global conflicts, none have exhibited the multifaceted complexities and rapid dynamics that will characterize such a conflict.

Just as war seems to be evolving toward something unprecedented, the rise of artificial intelligence (AI) in military operations has the potential to present an even more unprecedented evolution. This technology is advancing at an impressive rate, and its unpredictable development brings profound implications for the future of warfare. Although the U.S. military has applied AI to specific tasks—such as target recognition and autonomous navigation—its applications are on the verge of dramatic expansion. As AI systems evolve, they will increasingly permeate various operational aspects, from logistics and supply chain management to real-time tactical decision-making processes. It is not far-fetched to anticipate a future where AI becomes integral to military command and control systems, facilitating more rapid and comprehensive decision-making processes.

As we approach this transformation, it is evident that military organizations are not fully prepared for the disruptions that are likely to unfold.

This paper explores a reimagined approach to command and control that maximizes the benefits of advanced artificial intelligence technologies through innovative models and algorithms. This approach involves integrating AI into the command-and-control paradigm, potentially crafting a structure that balances the advantages of both centralization—ensuring coherent strategic direction—and decentralization—allowing for localized, adaptive responses to fluid operational conditions. This powerful fusion of traditional command structures with cutting-edge AI capabilities can facilitate a profound streamlining of global military operations, improving responsiveness and effectiveness.

Note: For better or worse, the term *command and control* has many expanded variants. For simplicity, I will follow Van Creveld's example and use the term *command* to encompass the many facets of command and control described in acronyms such as C3, C4ISR, C4I, etc.

Command + AI = Something New

War is an immensely complex part of the human experience, encompassing all the disciplines of human learning—from philosophy to physics—in pursuit of victory. The essential human task in war is command, and it includes everything from the design of operations in pursuit of a desired outcome to the coordinated application of violence at the local level. Effective command can overcome the human elements that limit military operations. Through command, humans cope with the forces of fog, friction, fear, uncertainty, and change. Also, through command, humans take advantage of the human qualities of imagination, courage, genius, empathy, and morality.

As we consider war in its totality, perhaps no other facet reflects the good and bad of humanity more than command.

But war extends far beyond human cognition to the physical world, and in the physical world, humans find it quite valuable to augment themselves with machines. Over time, these machines have become indispensable for winning wars. Humans fight wars with crossbows, catapults, battering rams, tanks, ships, airplanes, guns, and bombs. These machines displayed essential differences from their human operators as they were developed and used. Machines don't feel fear. Machines don't have empathy. Machines don't apply morality. Machines don't innovate. Machines don't get tired or distracted. Machines don't have a family or country to fight for. Machines don't think.

Over thousands of years of war, humans and machines have played very different roles--until now. This sharp distinction will change with the incorporation of artificial intelligence into military operations, which will blur the distinctions between man and machine in disruptive ways. These changes will be particularly evident in the process of commanding military operations.

Artificial intelligence has already changed operations by accomplishing tasks such as recognizing patterns—including identifying targets within clutter—and optimizing weapons to targets. In the near future, these use cases will expand to many more applications. In the short term, the machines will do what humans can already do, but better. Humans can analyze imagery, but machines can do it quicker and more accurately. Humans can write and read task orders, but machines can do this without getting tired or distracted. Humans can assess battlefield progress or lack thereof, but machines can do this more quickly and comprehensively without getting emotional about the results.

However, a turning point is coming when AI-enabled algorithms can do things that humans cannot because human cognition is limited, and established paradigms are difficult to overcome. For example, machines can spot connections between disparate data sets that often exist within separate parts of military operations (i.e., in the stovepipes of the J-staff or the individual intelligence agencies). Machines can optimize between hundreds of complex variables, whereas humans can only satisfy. Machines can recognize changes in conditions and respond long before these become apparent to humans. Importantly, machines are not inherently limited by paradigms and doctrines, where there is sizable evidence that humans are limited by their experiences and beliefs. When algorithms are trained to direct military operations at scale, it will be surprising if they do not identify novel strategies like Alpha-Go identified novel approaches to that ancient game.

It's also likely that machines will generate junk along the way, leading some to conclude that AI can never be allowed to participate in command. Learning to train AI in command will be non-linear and uncertain, allowing for strong opinions about AI on both sides.

Trust is the essential moderator in this process. As military forces incorporate AI into command, there will be opportunities to expand the level of trust humans place in AI. As trust is built through testing, evaluation, exercises, and experience, the role of machines in command will grow. At first, trust will come slowly. As our experience grows, so will our trust.

We are now asking ourselves: What role should machines play in command? As discussed above, a starting point in answering this question is to identify where machines can do the things humans can do but better. There are many discrete tasks within command where AI can be applied and evaluated quickly, simply because human experts can recognize when the machines are getting it right. Someone experienced in planning air refueling operations can quickly identify when an optimization algorithm is doing its job of matching aircraft to tankers. In cases like this, the evaluation and iteration process can go quickly, allowing trust to grow accordingly.

There will be roles where AI can do things that humans cannot, and trust is much more challenging to establish in these cases. For example, it is exceptionally difficult for humans to understand what is happening in a contested electromagnetic (EM) spectrum, as there is no accepted way for humans to visualize the competition within this part of the battlespace. Machines, in contrast, could abstract the EM spectrum into its mathematical components to evaluate the existing conditions and recommend actions. In this case, humans may be unable to follow the machines' process, or at least not follow it in a relevant timeline. It will be easier to trust machines when humans readily recognize if they are right or wrong and much more difficult when they cannot.

In addition, there are differences between live and virtual environments. In many cases, we can only develop warfighting algorithms in virtual environments because we cannot replicate the actual battlespace conditions for safety, security, cost, or logistical reasons. In the physical world, no test range can reproduce the full "dirty" spectrum expected in combat, a wave of hypersonic weapons, or a battle to protect and destroy satellites in low-earth orbit. Yet algorithms will likely be helpful in all these cases. We will be constrained to develop these in the virtual world and then deploy them to the physical world when needed. There will be a precarious period in battle when we will deploy an algorithm and run it "live" for the first time. There will be new stimuli in this live environment. It would be surprising if there were no unexpected developments when we deploy algorithms into the combat environment, but that is true for humans and machines today. The first time one goes into combat is different. We will always learn new things about ourselves and our equipment in the early periods of fighting. The novelty of the combat environment should not deter us from deploying algorithms into this environment. However, it does require an intentional approach.

This new approach will reflect changes in the human tasks of command. In the past, it has been necessary for the commander to prepare his mind for the required decisions in wartime and to prepare his staff and subordinate commanders to execute his intent. Commanders have used tools such as exercises and "rock drills" to do this. Likewise, preparing the command algorithms will be a crucial task for the commander—just as important as preparing the staff and subordinates. Since preparing these algorithms is not a singular event but a continuous, iterative process, success in wartime will require commanders to conduct continuous learning during preparation periods. This preparation cannot be shortchanged. Essentially, the commander, his staff, and the command algorithms are learning and developing together, training each other and creating the human-machine teaming practices that will carry over into battle.

As they do this, they will learn a potent division of labor between machines and humans. They will ask themselves: How can we combine the best of humans and machines to command forces in war? Through repeated exercises and simulations, they will likely learn new practices that create synergy together. Evidence suggests that the level of command effectiveness will elevate as humans and machines work together.

A vital part of this learning process will be asking the question: What is the proper role of humans in command? The answers may not be what we expect. The role of humans in human-machine command teams may be quite different than before. There are likely old tasks better accomplished by machines, old tasks shared with machines, and new tasks that humans must perform to help the machines.

As we begin this learning process, here are some hypotheses. Some will be right, others wrong (or proven to be more nuanced).

- The decision to go to war is a human decision.
- The general decision to engage with lethal force is a human decision. This decision includes the parameters of acceptable engagement, including *rules of engagement (ROE) that specify ge*ography, mission, timeframe, and parameters for engagement of targets.
- The identification of targets is a machine decision with human veto power.
- The selection of a strategy/tactic is a shared decision.
- The matching of targets with effects consistent with ROE and selected tactic is a machine decision with human veto power.
- The recognition of changing parameters is a machine decision.
- The suggestion of adaptive actions—including new strategies and tactics—is a machine decision.
- The adoption of adaptive actions is a shared decision.
- The decision to stop and disengage before the set timeframe is a human decision.

- The recognition of algorithms that are acting in ways inconsistent with desired behaviors is a human decision.

As the learning and preparation process unfolds, and repeated simulation allows for a deep appreciation of the operational problems, the command algorithms become an extension of the commander. Put another way, they are an essential element in the modern version of a *commander's intent*, much more reflective of the commander's true desires than a standard five-paragraph order. When this happens, these command algorithms enable a new, highly lethal form of *mission command*.

Mission Command Powered by Human-Machine Command Teams

In modern combat, we expect the overall commander to be unable to communicate with subordinate commanders consistently. The command algorithms, however, can be deployed to subordinate command nodes at the outset of the conflict. When this happens, the algorithms preserve connectedness with the commander's intent even when the commander is not physically connected. In addition, they do this at a much more granular level than a five-paragraph order. These command algorithms will do many of the tasks identified above quickly and accurately. The subordinate commanders and command nodes can rely on the command algorithms to help execute operations within the overall intent.

In our historical experience, the major weakness in mission command has been inadequate communication between the commander and subordinate commanders. The commander's intent has been expressed in general terms (i.e., five-paragraph format), which has put tremendous burdens on effective communication in the fog of combat. In many cases, the commander has been unable to express himself in written language, and subordinate commanders have not been fully able or willing to seek clarification. The result is more uncertainty in the most challenging of environments, preventing the full advantages of mission command. The communications gap is a part of the nature of war, and it has not changed.

In contrast, when the commander pushes command algorithms developed concurrently with the commander's evolving appreciation of the operational problem and trained through thousands of simulations, there is the potential to bridge the gap between commander and subordinates differently. Consistent with the hypotheses above, the subordinate commander and staff can share decision-making with the algorithms even when cut off from the overall commander. The algorithms offer a connectedness that can mitigate communication problems during decentralized operations.

But there is a trade-off. As identified earlier, algorithms deployed into the combat environment for the first time will likely experience anomalies. As the subordinate commander applies the commander's intent through command algorithms, someone will need to be able to recognize when the algorithms are deviating from the expected results. For example, the adversary may be protecting high-value targets through a sophisticated AI-spoofing technique, or the command algorithm may be expending exquisite weapons at too high a rate due to a deviation in how it calculates the probability of success. At this point, the subordinate commander and staff need to recognize the deviation first, which requires a deep appreciation for the commander's intent. Once the subordinate staff identifies the deviation, they need to be able to correct it without stopping operations. It will not be acceptable to revert to human-only command, as that will undoubtedly degrade the velocity of operations and place the unit in a worse situation. However, if we are not intentional about providing the proper training and tools, that is what will happen.

Being intentional about human-machine teaming in command means we must develop new tools and skill sets for humans as we develop and field command algorithms. Until now, the discussion has centered on humans being "in the loop," meaning that humans must affirmatively concur with each decision made, or "on the loop," where humans preserve the ability to stop decisions where they do not concur. In the reimagined world of command enabled by algorithms, neither describes humans' most critical role. Instead, we might say that, in healthy human-machine teams, humans are "inside the loop." They have been trained to recognize deviations from the intent and adjust the command algorithms to correct them as the algorithms continue to fight.

Even when the command algorithms are operating as designed, there will be a need for this "inside the loop" capability. There will be surprises on the battlefield that were not anticipated as the commander trained the algorithms in simulations. Some of these will be relatively small surprises—i.e., wartime modes of radars, novel decoys, etc. There will be some big surprises, too. The algorithms may recognize these surprises, adjust to them, and stay consistent with the commander's intent. Or they may not, which will require human intervention inside the loop.

A human-machine command system that allows for the centralized development and the decentralized execution of the commander's intent—while also allowing adjustments when deviating from this intent—would enable a powerful new approach to mission command. This system would preserve the advantages of centralized direction in military operations, including promulgating the commander's intent in a way much more representative of the actual commander. It would also take advantage of the ability of subordinate commanders to adapt to changing conditions and exercise initiative without waiting for approval. This system would be resilient to fracturing because it allows the "glue" of the commander's intent to be replicated down the chain of command and executed even when subordinates are cut off from the commander. In addition, it allows for battlefield adaptation at a swift pace while maintaining the critical connectedness between the commander and the subordinate commander.

This enhanced relationship between leader and subordinate has the potential to be replicated at all levels of command—from global to local. If the leader can develop the command algorithms through multiple simulations and iterations, the subordinates can execute the commander's intent quickly and faithfully, even in the fog and friction of combat. This approach could allow the Secretary of Defense to use command algorithms to allocate scarce forces and capabilities to combatant commands, the theater commander to set targeting priorities for the joint task force, and the local fires commander to match weapons to targets.

The (Abstract) Argument Thus Far...

IF

The commander and staff develop command algorithms through simulation and iteration,

AND

The subordinate commanders and staff develop an appreciation for the commander's intent,

AND

The command algorithms are deployed at the outset of combat,

AND

Subordinate commands have tools to adjust the algorithms while continuing the fight,

THEN

The subordinate commands can apply human-machine teaming to execute rapid and lethal combat operations consistent with commander's intent,

AND

When they recognize deviations from the commander's intent, they can operate "inside the loop" to adjust the command algorithms while fighting.

THE RESULT WILL BE

A new and powerful form of mission command replicable throughout all command levels.

A (More) Concrete Example to Illustrate the Argument

Given a potential scenario of defending Taiwan against a Chinese invasion…

In preparation:

- The National Security Council (NSC) runs a wargame to identify key decision points for the President of the United States (POTUS), including options for deterrence and escalation.
- Secretary of Defense (SD) develops critical decision points on the DoD level and assigns staff to build command algorithms to aid in these decisions. This begins a cycle of simulation and iteration to develop algorithms for global force allocation, force generation, and deployment. The SD regularly works with the staff to tune the algorithms and iterate on "what if" scenarios.
- The regional Combatant Commanders develop command algorithms to "set the theater," including communications planning, traditional and non-traditional logistics planning (including priorities for pre-positioning), force posture and beddown, and initial rules of engagement.
- The functional Combatant Commanders develop command algorithms to address operational problems within their functions given potential forces available and rules of engagement.
- Task force commanders and staff are activated to develop command algorithms to solve their assigned operational problems. For example, the Combined Joint Task Force – Taiwan (CJTF-T) is responsible for working with Taiwan in its defense. In simulation, these task forces develop an appreciation of the operational problem and a commander's intent to address it. In the case of CJTF-T, the commander and staff develop and tune algorithms to prioritize targets within the operating area, optimize fires against these targets, allow for collaboration with Taiwan's Defense Forces and that of allies, and create demand signals for deployed forces to "donate" fires in Taiwan's defense. The task force should develop command algorithms in concert with Taiwan, Australia, and Japan as much as possible.
- Combat units develop and tune command algorithms that connect to intelligence feeds and/or common operating pictures to characterize their operating areas. They incorporate detailed mapping of terrain, infrastructure, and pre-positioned logistics. During exercises, they develop the command algorithms to help them defend against threats and position themselves to provide theater fires for the task force.

As a crisis unfolds…

- POTUS works with the NSC to make the decision to go to war. Additionally, POTUS decides on objectives, authorities, constraints, and restraints.

- The SD and staff input POTUS decisions into the global command algorithms. SD staff initiates the algorithms that generate forces and deployment orders. They monitor outputs for deviations from SD intent and correct them as necessary.
  - *Example deviation:* Units receive deployment orders that are not executable due to real-world limitations.
- The regional Combatant Commanders receive their initial force allocation. Their command algorithms generate the optimal beddown of forces to support theater objectives.
  - *Example deviation:* Units are directed to non-functional bases suffering from cyber attacks.
- The functional Combatant Commanders receive their force allocation. Their command algorithms generate options for achieving the commanders' objectives within their functions.
  - *Example deviation:* The allocated forces are much less than simulated, and the command algorithm distributes capabilities over many objectives rather than concentrating on the most important priorities.
- Task force commanders receive allocated forces and beddown plan. The command algorithms begin planning operations, beginning with forces available immediately.
  - *Example deviation:* The common operating picture displays an order of magnitude more targets than expected.
- CJTF-T deploys its command algorithms to subordinate command nodes. These command nodes are deployed to survivable sites (such as a hide site in the first island chain). If the CJTF-T headquarters is cut off, each subordinate command node is prepared to direct fires in support of Taiwan's defense according to a predetermined order of operations.
  - *Example deviation:* Two subordinate command nodes think they have assumed primary responsibility. This is corrected when a deployed unit receives two different calls for fires, which are automatically rejected by that unit's command algorithm. The separate command nodes receive the rejection, which signals them to revert to primary and secondary.
- Deployed units begin combat operations. Their command algorithms respond to calls for theater fires by creating movement plans that position their launchers to be ready to fire in the timed window. These algorithms also help deployed commanders and staffs to execute tactics to evade detection and take advantage of prepositioned supplies and munitions. These algorithms report the unit's status and generate calls for support, including logistics support, as needed.
  - *Example deviation:* The command algorithm may not have access to accurate data about infrastructure and prepositioned logistics sites that have been attacked.

At every level of command, there is a powerful and practical connectedness through the command algorithms. Also, at every level, there is the potential for the algorithms to deviate from what is intended. Therefore, the challenge of human-machine teaming is to maximize the good while mitigating the bad. Commanders and their staffs maximize the good of command algorithms through their personal involvement in developing, training, and tuning these algorithms. Subordinate commanders and their staffs mitigate the potential of deviations by identifying when the command algorithms are deviating and using "inside the loop" tools to adjust them.

With this approach, command is a shared responsibility between humans and machines. This is a fundamental change for military forces, and it has important implications.

Implications for This Reimagined Approach to Command and Control

A commander's preparation for combat will be different. The development, training, and tuning of command algorithms will be essential as commanders ready themselves and their units for combat. In addition to training and exercises, commanders and their staffs must extensively engage in simulation and modeling. The foundation for successful human-machine teaming in combat will be built during the iteration required to prepare the command algorithms for deployment. Commanders who fail to engage in this process will deprive their units of this foundation.

Subordinate commanders and their staff must understand intent at a deeper level. As they must recognize when the command algorithms deviate from the original intent, they must understand how operations should be executed and compare that understanding with how operations are going in the real world. This will likely require a much more granular feel for the commander's intent, which is challenging to attain. Just as commanders must devote more time and effort to develop their intent as expressed in the command algorithms, subordinate commanders and their staff must devote time and effort to understanding this intent.

Subordinate commanders and their staffs will need the tools and training to operate inside the loop. The temptation will be to operate the command algorithms with an on/off switch...if these algorithms are not consistent with expectations, then you can turn them off and operate in "human-only" mode, analogous to a pilot who turns off the autopilot and flies manually. This is unlikely to turn out well. A better approach would be to give the humans tools and training that allow for adjustment while continuing the fight. Independent evaluation algorithms may become a tool that helps humans to diagnose what is happening. In addition, it may be possible for humans to manually control some inputs into the command algorithms without taking these offline.

Human-machine command offers potential solutions to modern challenges of command and control. As the character of war evolves in speed, scale, and complexity while maintaining the nature of fog, friction, and surprise, this human-machine command concept can help preserve effective command and control by leveraging the best of machines and the unique qualities of humans. The result is likely to outpace other approaches and offer significant advantages.

Human-machine command can be a powerful catalyst for adaptation in combat. In a previous SCSP paper, Major General (ret.) Mick Ryan and I argued in the Future Character of Conflict 2030:

> An organization's adaptive capacity—its ability to adapt on the battlefield and as an institution—has always been and will remain critical to success in war. However, the pace of war and the speed of change in the geostrategic environment mean that military institutions' existing approaches to learning, adaptation, and organizational decision-making about change must be reinvented. Victory will smile on organizations that "supercharge" their learning and adaptation capacity.

The human-machine command concept described in this paper is one way for a military force to supercharge adaptation. With access to the right inputs, the command algorithms will recognize when operations are not meeting objective criteria and make humans aware of this. The machines can suggest operational adjustments, and/or the humans can adjust the algorithms. This flexibility will be critical in fighting the battle of adaptation.

Human-machine command could streamline bottlenecks that exist in the current command structure. The United States has a significant bottleneck in the wartime chain of command as it goes from the Secretary of Defense (and a relatively small Joint Staff) to the eleven Combatant Commanders. In a crisis, it is difficult to see how the Secretary of Defense will be able to cope with all the requests from the Combatant Commanders. However, the Secretary of Defense could build a human-machine command system to streamline this chokepoint with the help of command algorithms that have been tuned to the Secretary's logic and tested in multiple simulations.