

Accelerating Decision-Making Through Human-Machine Teaming

1st Lt Ian Palmer

DISCLAIMER: The opinions expressed in this essay are those of the author and do not necessarily reflect the official policy or position of the Department of Defense or any other U.S. Government agencies.

The United States finds its military advantage waning. As the US dedicated twenty years to the Global War on Terror, its strategic competitors undertook a deliberate campaign to develop military capabilities that counter the technological and doctrinal strengths of the US military. Recognizing the need for a new effort to regain a military advantage, the “Third Offset” began in 2014 with the Defense Innovation Initiative.¹ There are five key technologies in the Third Offset: learning machines, human-machine collaboration, assisted human operations, human-machine combat teaming, and network-enabled autonomous weapons. Even as these technologies reach maturity and are delivered to the Air Force, though, there currently exists no doctrinal framework to employ them at the operational level of warfare. The Air Force must create new doctrinal mechanisms that support Third Offset technologies. This paper first explores several human-machine teaming frameworks, then advocates for their incorporation into existing doctrinal authorities. It also explores the challenges associated with those changes and the importance of undertaking change despite the risk.

The First and Second Offsets

This is not the first or even second time the US has faced this challenge. In the 1950s, the United States and its European allies were competing against the Soviet Union, which fielded a significantly larger military. The “First Offset” strategy employed by the United States featured an emphasis on nuclear weapons, delivery vehicles, and a doctrine of nuclear deterrence.² This preserved peace in Europe until the 1970s, at which point the Soviet Union equaled or even surpassed the United States in nuclear weapon quantity and effectiveness. In response U.S. policymakers initiated the Second Offset strategy. This effort, spearheaded by the Defense Advanced Research Projects Agency (DARPA), fueled the development of the precision-guided munitions, low-observable aircraft, and long-range active radar-guided missiles that remain the key enablers of Air Force airpower today.³ Getting the Third Offset right, as Deputy Secretary of Defense Bob Work argued, will “maintain and perhaps advance the competitive advantage of America and its military allies”⁴, just as the first two offset strategies did.

Artificial Intelligence

Artificial intelligence (AI) is “the study of the computations that make it possible to perceive, reason, and act.”⁵ AI systems excel at processing large volumes, varieties, and velocities of data –

¹ Secretary of Defense Chuck Hagel. “Reagan National Defense Forum Keynote.” Speech, November 15, 2014.

² Peter Grier, “The First Offset”, *Air Force Magazine*, June 2016, 1-2.

³ Rebecca Grant, “The Second Offset”, *Air and Space Forces Magazine*, June 24, 2016, 2-8.

⁴ Deputy Secretary of Defense Bob Work. “The Third U.S. Offset Strategy and its Implications for Partners and Allies.” Speech, January 28, 2015.

⁵ Patrick Winston, *Artificial Intelligence: Third Edition* (Reading: Addison-Wesley, 1992), 5

the ‘3 Vs’ of big data.⁶ As computing infrastructure and networks expand, AI systems can ingest more and more data and provide greater advantages over humans performing the same task. IBM’s DeepBlue became the first machine algorithm to defeat the world chess champion in 1997,⁷ and in 2015 AlphaGo defeated the world Go champion.⁸ These leaps were enabled by large-scale computing and data innovations.

Modern machine learning algorithms, though, are often brittle and fail to adapt to new settings. When the input data to an algorithm differs significantly from the data originally used to train it, these systems can fail in unexpected ways.⁹ Adversaries, too, can exploit learning systems to encourage undesired behaviors.

Artificial intelligence, then, is a powerful (but often misunderstood) tool to create systems that can carry out tasks autonomously. Often, though, AI systems are used in conjunction with high-level human supervision, a practice called human-machine teaming.

Human-Machine Teaming

As previously discussed, machine algorithms excel at processing the ‘3 Vs’ of data but can struggle to generalize beyond the limits of the data used during training. Humans are naturally capable of making decisions in high-uncertainty or novel situations but struggle with processing high volumes, varieties, and velocities of data – exactly the opposite. This is Moravec’s Paradox: machine algorithms find tasks that are hard for humans easy and vice-versa.¹⁰ Human-machine teaming, then, seeks to take advantage of the strengths of humans and machines while minimizing their weaknesses. There are three general frameworks for human-machine teaming, each with a different level of machine autonomy:¹¹

Human-in-the-loop processes give human decision-makers the ultimate approval authority in an automated process. No actions are performed without the explicit approval of a human. This mode is most useful when oversight of a process is critical, but it reduces the usefulness of having an autonomous system in the first place.

In **human-on-the-loop** processes, humans have a ‘veto’ authority that can override an impending action that will otherwise occur. This framework balances autonomy and supervision, but it can be challenging to determine exactly how or when a human can intervene.

⁶ Dr. Peter Layton, *Algorithmic Warfare: Applying Artificial Intelligence to Warfighting* (Canberra: Air Power Development Centre, 2018), 10.

⁷ Andrew McAfee, “Did Garry Kasparov Stumble Into a New Business Process Model?”, *Harvard Business Review*, February 18, 2010, 1-4.

⁸ Tanguy Chouinard, “The Go Files: AI computer clinches victory against Go champion”, *Nature* (March 2016).

⁹ Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (Upper Saddle River: Prentice Hall, 2010), 409-410.

¹⁰ Layton, *Algorithmic Warfare*, 58.

¹¹ *Ibid.*, 28.

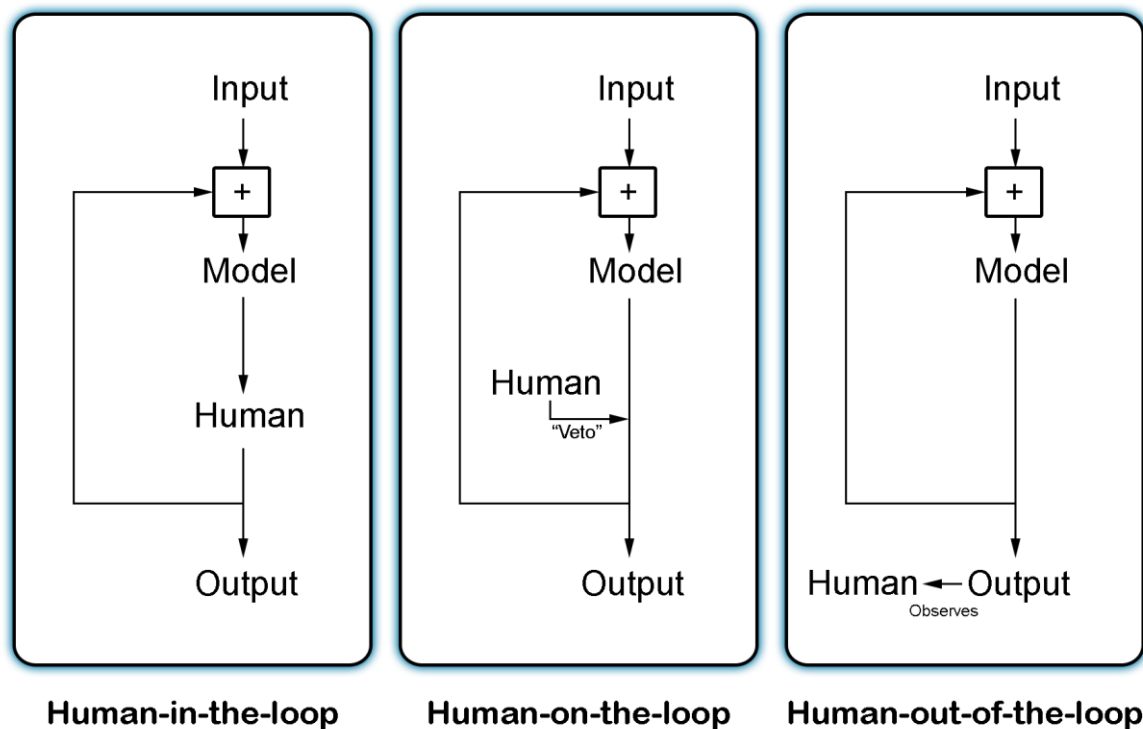


Figure 1: The three basic frameworks of human-machine teaming. Source: 1Lt Palmer

Human-out-of-the-loop differs from the other two frameworks in that there is no human involvement in the process. The machine makes all decisions independently.

The three frameworks are summarized in Figure 1. There is no ‘best’ framework; each has benefits and drawbacks that make it applicable to certain situations.

“Do Better Things”

The Air Force might be tempted to integrate artificial intelligence into its operations by simply using these new capabilities within the confines of existing doctrine. The appeal of this approach is its simplicity: Airmen already familiar with USAF doctrine would not need to learn any new processes or systems in order to leverage the increased information processing power of AI. We would simply do the things we already do, but slightly more effectively. Dr. Peter Layton describes this as the “do things better” approach to algorithmic warfare.¹²

That approach is not sufficient to maintain a USAF warfighting edge. China and Russia, the two primary strategic competitors to the United States, have studied US warfighting doctrine and have developed technologies and doctrines designed specifically to counter its strengths. The alternative is the “do better things” approach, which emphasizes the use of semi-autonomous to autonomous machines in active roles across the battlespace. Machines would no longer simply be situational awareness-enhancing tools, but rather active participants working collaboratively with humans.

¹² Layton, *Algorithmic Warfare*, 36-39.

These machines could, as Layton puts it, “wage war semi-independently, only needing human guidance from afar occasionally.”¹³

The key difference between the two is the relative duration of the Observe-Orient-Decide-Act (OODA) loop occurring at the operational level of warfare. In traditional doctrine, the time it takes for a human to make a decision will always be the limiting factor. Acceleration through the integration of AI battle management tools will only speed up the process to a certain point. To gain a decisive advantage over an enemy, though, the Air Force must seek further improvements by reducing the role of humans in the decision-making process whenever possible.

Teaming in Practice

The OODA loop of modern USAF doctrine can be illustrated through the Defensive Counterair (DCA) mission. In DCA, rapid decision-making is critical in order to ensure the survival of friendly forces.¹⁴ Three decision-makers in particular play significant roles: the ID Authority, Commit Authority, and Engagement Authority.¹⁵ All three of these are key components of the OODA loop that must be completed before an engagement can take place. With humans in the loop, however, there are vulnerabilities that adversaries may exploit to delay the process and corresponding advantages to the use of human-machine teaming.

As an example of DCA in practice, imagine that the Control and Reporting Center (CRC) detects a suspect unmanned aerial vehicle (UAV) using its acquisition radar. The CRC disseminates this information to the Air Operations Center (AOC). At the AOC, the Joint Forces Air Component Commander (JFACC), acting as the AADC, holds engagement authority and delegable commit authority. Meanwhile, the identification authority uses the criteria established by the AADC to achieve CID. Once CID is met, a capable asset commits against the track, and engagement authority is given, the hostile UAV may be destroyed.¹⁶

This construct works well enough against one threat at a time. An adversary looking to exploit this system, though, might try to overwhelm the OODA loop by flying several UAVs at the same time in different geographical areas of the theater. This will challenge the decision-making authorities, who now have to keep track of several threats in different places and advance several kill chains simultaneously.

There is a partial solution to this within existing doctrine: the regional air defense commander (RADC) or sector air defense commander (SADC).¹⁷ By planning to “divide and conquer”, the AOC and RADC/SADC can increase their capacity to track and engage threats across the theater. However, it is easier than ever for adversaries to exploit the principle of mass and induce even further strain on the DCA OODA loop by flying more and more UAVs across the theater as these systems become more capable and less expensive. There is no mechanism for a RADC or SADC to further delegate their authorities even if the amount of information overwhelms their information processing bandwidth. Even if there were, that would only push the problem down to

¹³ Layton, *Algorithmic Warfare*, 40-46.

¹⁴ Joint Doctrine, JP 3-01, *Joint Countering Air and Missile Threats*, I-7.

¹⁵ United States Air Force Doctrine, AFDP 3-01, *Counterair Operations*, 11-12.

¹⁶ AFDP 3-01, 20-24.

¹⁷ JP 3-01, II-14 and AFDP 3-01, 11.

the next human in the decision-making chain. As long as humans are making the decisions, they will be the limiting factor in the OODA loop.

Instead, Air Force doctrine must take inspiration from existing human-machine teaming methodologies and create a framework within which human operational control elements (e.g. the AADC) can delegate a varying level of authority to machines based on operational requirements. There are a few requirements to such a framework: importantly, it must be flexible enough to accommodate the range of military operations. The Joint Forces Commander (JFC) and JFACC must have the ability to adjust the mode of human-machine teaming (whether humans are in, on, or out of the loop) at any time. It must also be general enough to adapt to changes in technology.

Imagine instead that Air Force doctrine defined the three human-machine teaming methods above and established mechanisms by which operational authorities could explicitly select one of those modes to guide their own decision-making process. The selection of a teaming mode could be disseminated via existing mechanisms like theater Special Instructions (SPINS) and change at any time depending on operational requirements. In practice, this change should start with defining human-machine teaming and the three modes in AFDP 3-30, *Command and Control*. Specific domain doctrine, like AFDP 3-01 *Counterair*, AFDP 3-03 *Counterland*, and AFDP 3-12 *Cyberspace Operations*, should be modified to reflect that the authorities related to operational employment (i.e. the ID Authority in counterair operations) should be human-machine teams. Just as in the traditional delegation of authority, though, it should be made clear that humans still bear ultimate responsibility for their assigned duties.

As in the last scenario presented above, consider a swarm of unmanned aerial vehicles attacking simultaneously across a JFC's geographic area of responsibility. The JFACC, either in real-time as a response to the threat or ahead of time upon receiving intelligence that an attack was imminent, decides that human-on-the-loop teaming is the most appropriate mode for his or her engagement authority. That decision can be communicated through existing channels: the SPINS if the change is made prior to the attack, or simply voice if it is made in real-time. By implementing human-on-the-loop teaming, the JFACC prevents the theater OODA loop from being bottlenecked by a single decision-maker forced to authorize every engagement. It also prevents the JFACC from being removed from the loop entirely, as might happen if the JFACC were to simply provide a blanket authorization to engage every UAV. Instead, the JFACC can monitor a capable machine as it processes data faster than any human could. In the event that the machine made a decision that did not meet the JFACC's intent, he or she could override it. In that way, engagement decisions in line with the JFACC's intent could be made quickly, maximizing the time that assets have to complete their engagements and defeating a threat that could have otherwise overwhelmed theater defenses.

This concept is applicable beyond just the air domain. AFDP 3-12, *Cyberspace Operations*, outlines the role of the 616th OC and the authority to conduct defensive cyberspace operations (DCO) exercised by CDRAFCYBER¹⁸, which are roughly analogous to the role of an AOC and the engagement authority held by a JFACC. Many of the same concepts apply: today's cyber domain is becoming increasingly contested as threat actors increase in number and begin to use

¹⁸ United States Air Force Doctrine, AFDP 3-12, *Cyberspace Operations*, 17-19.

AI themselves. As AFDP 3-12 notes, timeliness is a critical factor in DCO.¹⁹ Completing the OODA loop and executing an appropriate defensive measure increases the probability that DCO can prevent data loss, malware attacks, or even physical damage to Air Force systems. CDRAFCYBER, as the authority to employ DCO, should determine the most appropriate mode of human-machine teaming based on the expected threat and the capability of the machine making decisions. Just as in air operations, this will maximize the probability of success while ensuring the commander's intent is met at all times.

There are several factors that decision-makers should consider in the selection of a human-machine teaming mode. The examples above use the quantity, capabilities, and locations of threats as the primary factors driving decisions; these considerations can be grouped together as operational requirements. The capability of the fielded battle management systems is another critical factor given the rate at which AI and machine learning technologies are developing. Additionally, political and strategic considerations may be a part of the decision. In an environment with restrictive ROEs, humans should retain more oversight of the process. Lastly, AFDP-1 notes several situations that have historically required a level of centralized execution, including strategic attack, covert/ clandestine operations, and high-value targets.²⁰ In those situations, the relevant commanders and authorities should remain in the loop in order to maximize human oversight. Together, all of these factors are summarized in Figure 2.

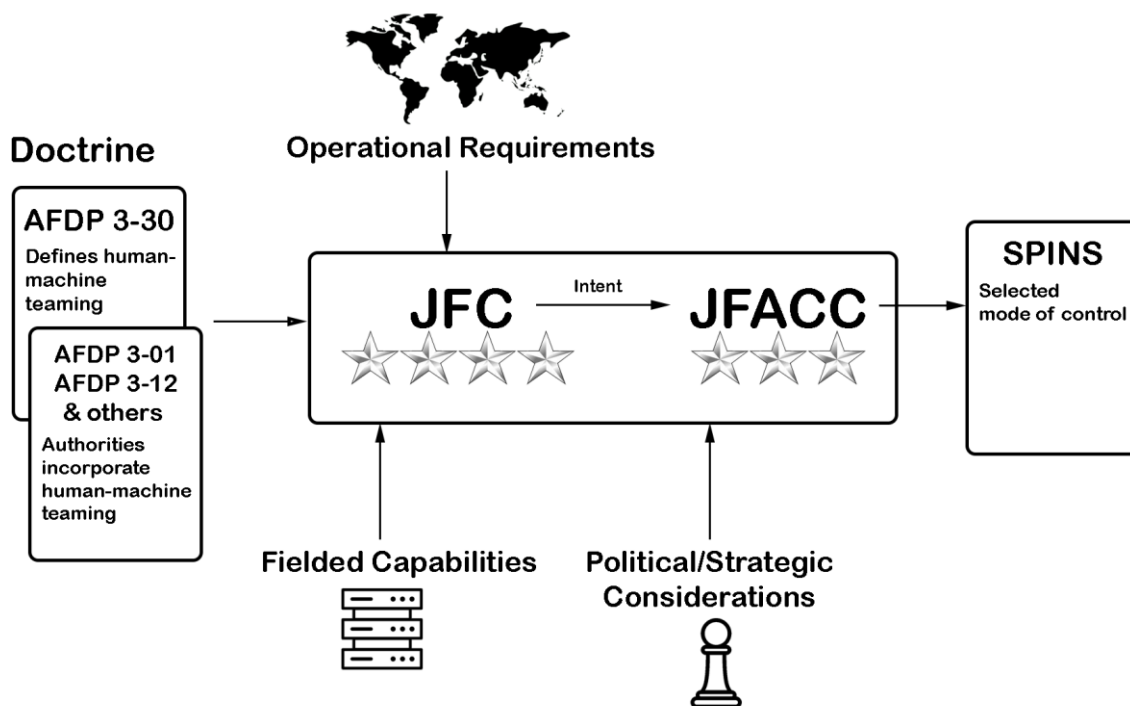


Figure 2: The process by which the JFC and JFACC select a mode of human-machine teaming.
Source: Clip art world map and 1Lt Palmer

¹⁹AFDP 3-12, 25.

²⁰ United States Air Force, AFDP-1, *The Air Force*, 14.

Challenges

Delegating any amount of authority to an artificial intelligence-based system is an extremely discomfoting idea for any military decision-maker, and rightfully so. AI systems behave as “black boxes” and have little ability to explain the reasoning behind their decisions. Centuries of military doctrinal development focused on the relationship between command and subordinate elements is being disrupted. Leaders at all levels will have to buy in to the change and fight through challenges that arise in order to make such a change possible.

AI systems can also make mistakes. When an AI system inevitably makes a decision that humans view as a mistake, how readily will humans be able to trust that system again? AI in any critical application, from self-driving vehicles to battle management software, faces this challenge. Humans may stop using the tool entirely after a single error, even if it objectively surpasses human performance on average.

It’s clear that the changes proposed here will not be easy to implement. The greatest risk, though, may be inaction. With the United States’ strategic competitors aggressively pursuing AI technologies and integrating them into their militaries, the Air Force could very quickly find itself at a tactical and operational disadvantage against a force using AI to its full potential. Implementing human-machine teaming in operational doctrine is not the entire solution for the USAF, but it is a viable first step and the potential benefits outweigh the risks.

The Path Forward

This is not the first time that the introduction of AI has disrupted a field traditionally dominated by human thinking. In 2005, the website Playchess.com hosted a chess tournament in which human and machine teams were permitted to collaborate in any way they saw fit.²¹ Some teams used only computer input, while other teams featured Grandmasters and state-of-the-art supercomputers working side-by-side. Across the board, human-machine pairings outperformed individual humans or computers alone. The overall winners, however, were surprising. They weren't Grandmasters and didn't have the most advanced chess algorithms; they were average chess players with average computer counterparts. Their advantage came from their use of superior human-machine teaming strategies to benefit from both human intuition and machine analytics, allowing them to defeat adversaries far stronger than themselves.²²

Implementing human-machine teaming into Air Force doctrine is, at its core, the same problem. Improving and modernizing doctrine, then, is of paramount importance to ensuring the United States’ continued security in the age of great power competition. The proposals here – to write human-machine teaming into doctrine, use those teams in existing operational authorities, and give commanders the ability to change the mode of teaming as required – are the best first steps for the Air Force to take on this challenge. Ultimately, change is difficult – but the alternative is failure.

²¹ Layton, *Algorithmic Warfare*, 26-27.

²² Garry Kasparov summarized that "weak human + machine + better process was superior to a strong computer alone and, more remarkably, superior to a strong human + machine + inferior process."