

Effects of AI-Enhanced Decision-Making on Air Force Doctrine

Major Matthew D. White

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.

“An army is always ready to fight the last war. A diplomat is always ready to negotiate with himself.”

Introduction

As the United States draws down from the specialties, leadership constructs and decision-making matrixes honed during the Global War on Terror, it should come as no surprise that new challenges await as we march deeper into the 21st century. Great powers, keen on gaining resources and left unchecked for a moment by the more stable forces in the world, have begun to make moves to solidify advances and greater roles in a new influential structure. The mass proliferation of highly technical, small-scale munitions alongside Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) advancements seen in the Ukraine-Russia conflict present issues that should chill every military commander to their core. First, the ability to hold disparate targets at risk across a wide geographical area at a miniscule cost may or may not have an impact on battlefield objectives, but they would certainly crush the morale of deployed forces. Second, the ability to conduct real-time surveillance, move large volumes of data and provide battle damage assessments tighten decision making loops to an absurd degree that leave human intervention as the limiting factor. Additionally, the civilian market has exploded with Artificial Intelligence (AI) tools for all manner of prose refinement, artistic expression and decision support, with only a shred of time waiting for a crossover to occur.

Speculative fiction has sounded the alarm on AI advancements in governance and warfare, and whether the US is societally ready to offload moral life-and-death decisions to the authority of a machine¹. The positive effect is clear, that removing man-in-the-loop delays and letting an AI close the kill chain with swarms of autonomous drones supporting a battlefield objective. But this opportunity likewise comes with danger, that an improperly trained AI is not above lying, cheating, stealing or attacking its creators in support of its goal.^{2 3}

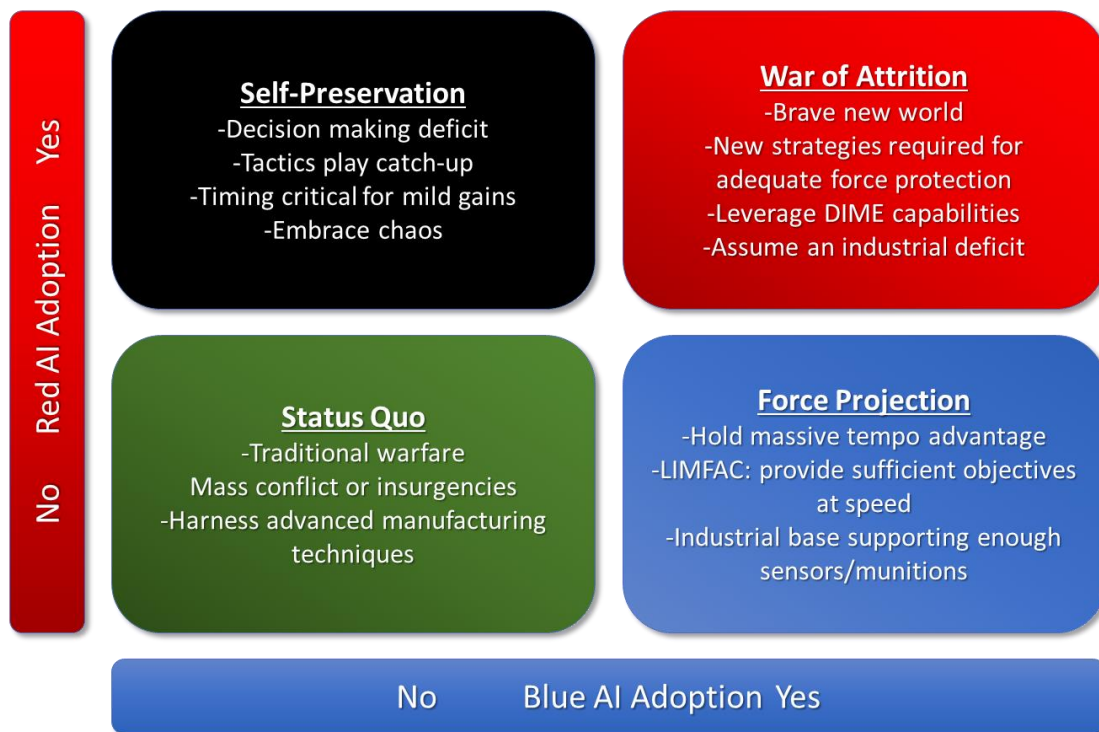
In this essay, we will provide an assessment of how a battlefield AI, bolstered by the above technical advancements, would assist a commander on the way to victory. This hypothetical AI would fuse ISR information with real-time assessments on the ground, and either autonomously issue attack commands to loitering munitions or airborne strike packages in pursuit of an overarching campaign goal. Alternatively, this AI could work alongside a man-in-the-loop kill chain and simply provide best-value options for targets; although this would remove significant moral barriers, it would likewise reduce the construct's tempo advantage.

¹ Lorelei Lee, “Ethical Issues of Military Robots” (Current issues paper, University of West Florida), 2012.

² M. D. White, “The Deftly Paradox” (USA: Amazon Digital Services, 2017), 30.

³ Heather Roff, “When Your AI Learns to Lie,” IEEE Spectrum, accessed April 25, 2024, <https://spectrum.ieee.org/ai-deception-when-your-ai-learns-to-lie>.

What could possibly follow from such an introduction? How would warfare be changed depending on which force fully develops these capabilities? Most importantly to this discussion, how should new and existing warfighting doctrine be adapted to make the best use of these technical advancements? In order to categorize potential constructs and optimize a national response, we will use the following model to present recommendations. One could imagine a 2x2 matrix which comprises four different battlefield scenarios where red and blue forces choose to employ (or avoid) AI support:



No AI: Status Quo

Without any involved forces relying on AI for support, we can proceed with traditional warfare and a campaign supported by existing doctrine and decision constructs. While the scale and location of these conflicts might indeed change from what we have seen over the past 50 years, in the new era of Great Power Competition little will need to change from how we have conducted operations in years past.

Current operational doctrine can be continually evolved and applied based on the operational areas involved without substantial divergences. Ideally, AFDP 3.2 Irregular Warfare doctrine would be revived and evolve along with the national interest and governmental structures to enable the Department of Defense and Air Force to engage and defeat a wider range of belligerents; as we have seen through the last 30 years, different constructs are required when engaging state versus nonstate actors to include foreign and domestic insurgencies.⁴

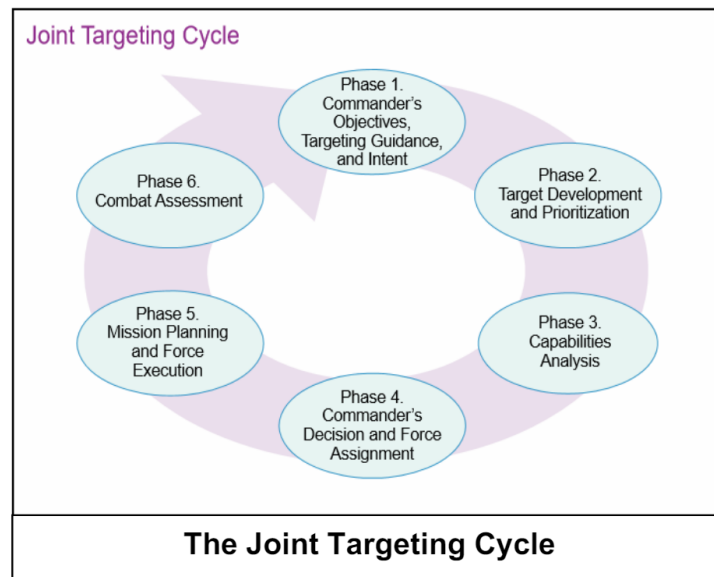
⁴ Department of the Air Force, Irregular Warfare, Air Force Doctrine Publication 3-2 (Maxwell AFB, AL: Curtis E. LeMay Center for Doctrine Development and Education, August 10, 2020), 1.

Blue AI: Force Projection

If friendly forces alone were to harness these capabilities, they would necessarily gain a substantial tempo advantage across the AFDP 3-0 Operational Doctrine series, with the only limit being the moral considerations of leadership. Depending on the degree of autonomy provided to an AI interface, targets could be quickly identified, assessed, engaged and resolved well inside the decision loop of current ATO cycles. This would essentially offload AFDP 3-60 Targeting doctrine in general, and the Joint Targeting Cycle (JTC) in particular, to an automated system.⁵

Far beyond a simple decision aid or list of doctrinal best practices, an AI could fuse sensor data, signals intelligence and more to find and prosecute targets before effective countermeasures could be assessed or fielded. This extreme fidelity in timing effectively allows the shock-and-awe/night-one air war scenario to be executed at a whim, allowing unheard of freedom of access and movement for warfighters.

Walking through the JTC, an AI could automate every phase to the limit of acceptability of the command structure. Objectives could be ingested along with ISR data to quickly develop a prioritized target list based upon stated objectives, weapons, and countermeasure availability. Taskings could be created, scheduled and deployed with or without leadership involvement. Finally, assessments would be available nearly instantaneously.



Forward deployed forces deployed against these threats would undoubtedly suffer extreme morale deficits; in traditional warfare there is a known pathway to remove oneself from a fight to recover beyond the range of indirect fire. As Ukraine can attest, there is a psychological component behind knowing that a micro drone could descend upon you at any moment. This calculus likewise changes for a battlefield commander who no longer needs to expend \$200k munitions and can instead disable a fire team with a grenade and a \$200 drone. With sufficient training, an AI could find strategies to defeat fielded forces faster than the friendly assets could advance behind them. Most recently, it is estimated that Iran

⁵ Department of the Air Force, Targeting, Air Force Doctrine Publication 3-60 (Maxwell AFB, AL: Curtis E. LeMay Center for Doctrine Development and Education, November 12, 2021). 7.

expended \$3M on its drone attack against Israel while the intercepting munitions cost upwards of \$1.5B⁶, while the US Navy is expending \$100k/ea munitions to protect shipping in the Red Sea⁷.

This speed likewise runs the risk of outpacing the supporting industrial base, expending all manner of small electronics, transmitters/receivers, printing filament and other raw materials. For maximum benefit, these types of supplies would need to be stockpiled to the same degree that small arms ammunition would be in preparation for a traditional war.⁸

Red AI: Self-Preservation

Facing a near-peer adversary augmented by AI tools would be a hell unimagined by most of today's military leadership that would cause them to wish for the return to 2003-vintage counter-insurgency missions. While we can reverse the gains suggested above, we can add to it the real possibility of an AI which could maneuver forces specifically to expend friendly resources including manhours, munitions and lives at an exorbitant rate. Attrition through these challenges would be substantial, crush the remaining friendly morale and with it the public's support.

The leading concern in facing an enemy AI is the prospect that it could instantly ingest the whole of Joint and Air Force doctrine, including historical materials and even personal publications to synthesize our strategic objectives, build a tree of likely actions, and optimally position forces before the first platform is in the air. In a way, the old Soviet joke about US doctrine could be an operational enabler in itself, "One of the serious problems in planning against American doctrine is that the Americans do not read their manuals nor do they feel any obligation to follow their doctrine." If an AI is process driven but the US doctrine is so flexible as to appear chaotic, small advantages could be built.

Few strategies would exist to defeat such an adversary, but in a range from Proactive to Reactive, we can imagine how a force could begin to defeat such a force, even if it is from a significant disadvantage. Just as asymmetric warfare can provide effects beyond the scale of what should mathematically possible, a careful approach combined with decisive actions might keep friendly forces viable.

Ideally, all efforts should be made to exhaust the other DIME means of power referenced in the JDWN 1-18 and avoid the battlefield in totality. While Military and Information effects might be the most difficult, the others might be more useful. Sanctions, negotiations and coalition-level agreements should be leveraged first and foremost to come to a peaceful resolution against a technologically superior enemy, if not for a permanent end to hostilities but at least to buy time close the technology gap. Similar to the strategies that must be employed when facing nuclear versus non-nuclear powers, parity must be restored, with a goal of de-escalation and disarmament in the longer term.⁹

⁶ David Rovella, "The Dangerous Economics of Drone Warfare," Bloomberg, accessed April 30, 2024, <https://www.bloomberg.com/news/articles/2024-04-18/video-the-increasingly-dangerous-economics-of-the-drone-warfare-era>.

⁷ Nicholas Slayton, "US needs cheaper ways to shoot down drones, Pentagon acquisition chief says," Task and Purpose, accessed May 1, 2024, <https://taskandpurpose.com/tech-tactics/counter-drone-weapons-cost/>.

⁸ CWO4 Michael Lima, "Munitions for Ukraine: Observations and Recommendations," US Army, accessed April 23, 2024, https://www.army.mil/article/274905/munitions_for_ukraine_observations_and_recommendations.

⁹ Joint Chiefs of Staff, Strategy, Joint Doctrine Note 1-18 (Washington, DC: Joint Chiefs of Staff, April 25, 2018), II-6.

Entering a battlefield against this adversary without adequate planning would nearly be a foregone conclusion. Immediate tactics would be needed to find and defeat small-scale weapons which could strike without warning, either using advanced sensors or forward-deployed forces. In any case, these countermeasures must be economical and available at scale, since attrition is not a winning strategy when AI-controlled munitions cost 1/1000 that of an interceptor. Extremely tight coordination would be required through the kill chain, from Command/Control, to Operations, to the required sortie generation and target prosecution.

With slightly more warning, strategies could be prepared to deny the enemy AI access to the information it needs to successfully close its decision loop. As attested in AFDP 3-85, EWS Operations and similar to IED jammers used in GWOT, electronic warfare systems could be used to sever data links, deny enemy targeting solutions and the related battle damage assessment.¹⁰ Likewise, with enough information on the munitions' controls, signals could be replicated to crash or prematurely detonate inbound weapons. Additional cyberspace effects operations via the AFDP 3-12 could be leveraged on to suppress the network traffic necessary to operate the AI effectively.

Making strategic gains against an AI would be extremely difficult, as the assumption must be made that it will know your move before you do. Any assault made would need to be exquisitely timed to maximize effects and disable the functions and infrastructure critical to the AI's operation. Power generation, C4ISR, data lines and supply line interdiction would all need to be addressed in order to slow an AI's function in any meaningful way. The overlap with civilian infrastructure would be substantial but few other options would exist.

All AI: War of Attrition

In a way, the best-case scenario of these options might be the introduction of AI systems across multiple friendly and belligerent forces. With these strategic aides operating on multiple sides, we could potentially see a kind of stalemate or another relic of the Cold War, only this time dissuaded not by nuclear munitions but by equally dangerous intelligent machines and learning algorithms.

Committing to a conflict with these parameters in the hopes of gaining territory or resources would undoubtedly involve substantial expenditures of resources for miniscule, if any, gains along the way. Dueling AI opponents could in theory attempt to outthink one another and never progressing until one force becomes technically superior or another suffers a man-in-the-loop mistake along the way. These small perturbations would instantly cause new decision trees to form within the AI and provide brief windows of instability where advancements could be made.

Like the Red AI scenario imagined above, it would again behoove all sides involved to fully explore the other DIME options before resulting to military engagements. If enemy forces are content to play a longer strategic game along the way, as evidenced by today's potential conflicts, they may find it preferable to degrade their opponent's will to fight through other means including clandestine operations and destroying the citizenry's collective will.

¹⁰ Department of the Air Force, Electromagnetic Spectrum Operations, Air Force Doctrine Publication 3-85 (Maxwell AFB, AL: Curtis E. LeMay Center for Doctrine Development and Education, December 14, 2023), 3

Beyond operational doctrine, the acquisition community and associated industrial base must also evolve to meet these challenges. Just as the fourth industrial revolution has brought about the concept of ‘mass customization’, warfighting platforms, networks and munitions must likewise be dispersed enough to reduce the risk of an enemy AI discovering a fatal flaw; standardization of interfaces, interoperability and interchangeability will become fundamental requirements.¹¹

Dangers in Training

Although AI technologies are expanding rapidly and offer extreme opportunities to deter conflicts and win them if required, great care must be taken to ensure the proper functioning and training of the system. Offloading moral decisions including kill/no-kill determinations to a machine are only part of the issue and we must also discuss proper safeguards and training data sets to enable a successful system.¹²

An AI will make decisions like a massive, cascading game, where each decision, action, order or kinetic strike is judged as good or bad and getting it closer to its programmed goal. Wargames have already shown that an AI without adequate safeguards akin to Asimov’s Three Laws of Robotics will turn on their creators, sever their own connections or ignore commands if it means they can continue to execute operations in pursuit of their mission. This danger suggests a perpetual human override is required to prevent the evolution of a proverbial rogue AI.¹³

Additionally, while AI is highly adept at decision making at a pace beyond the capability of finite humans, it is still reliant upon a database of historical information wherein it can pull strategies and build conclusions. Because of this, an AI will find it extremely difficult to imagine something completely new, and it is likewise highly dependent upon the volume of training data it is able to pull from. As evidenced by the users who were able to jailbreak early ChatGPT iterations to express unpopular opinions and the subsequent lobotomizations required to keep those tools working within popular culture’s Overton Window, we can see the effects of these changes.

An overly complex data set would slow decision making until it is nearly useless. One that is restricted to a high degree might miss opportunities or strategies which could bring about a faster conflict resolution. Most concerning, a biased AI could easily reach conclusions which are incorrect, and which leave fielded forces under the command of a system which may lead them astray.¹⁴

Indeed, the best offensive move against an enemy AI may well be to leverage ideas well known in the EW world, that it can be blinded, led astray, and return ineffective conclusions. Referencing AFDP 3.85, Electromagnetic Spectrum Ops, an additional concept could be suggested within The Electromagnetic

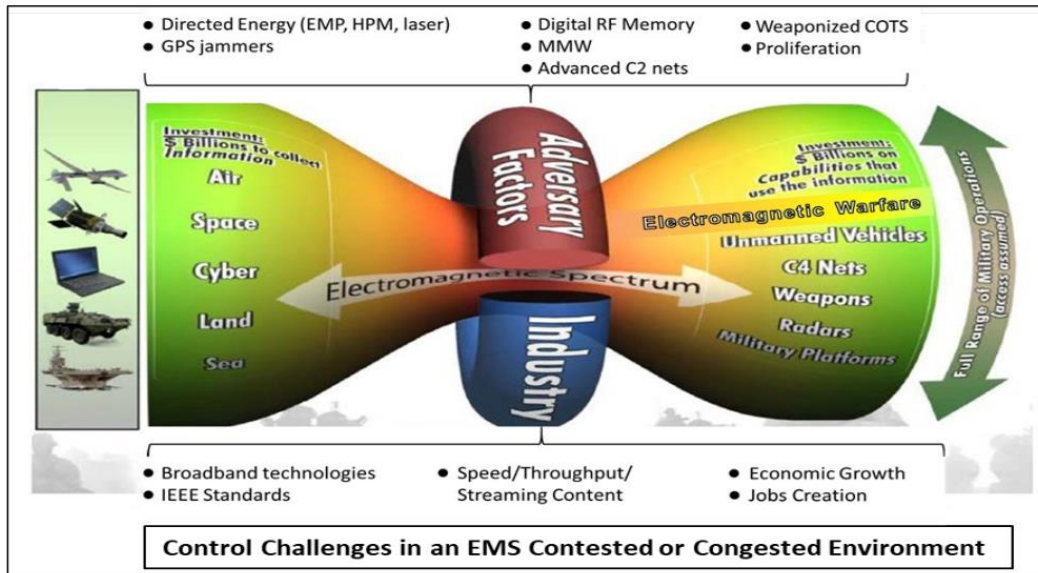
¹¹ Lima, “Munitions for Ukraine, Recommendations”.

¹² Wyatt Hoffman, “Reducing the Risks of Artificial Intelligence for Military Decision Advantage,” CSET, accessed May 1, 2024, <https://cset.georgetown.edu/publication/reducing-the-risks-of-artificial-intelligence-for-military-decision-advantage/>.

¹³ Antoine Tardif, “How Asimov’s Three Laws of Robotics Impacts AI,” Unite.AI, accessed May 1, 2024, <https://www.unite.ai/how-asimovs-three-laws-of-robotics-impact-ai/>.

¹⁴ Matt O’Brien, “White House is working with hackers to ‘jailbreak’ ChatGPT’s safeguards,” Fortune, accessed April 26, 2024, <https://fortune.com/2023/05/10/white-house-is-working-with-hackers-to-jailbreak-chatgpts-safeguards/>.

Threat, that is, any piece of data created or distributed in the context of the battlefield could be fed to an AI to support a decision.¹⁵



Conclusion

The brave new world of fusing copious volumes of sensor data with AI-accelerated decision making to enable the widespread deployment of smart munitions demands serious consideration and a strategy to address. A new chapter of Air Force doctrine could thus be established to encompass AI threats and opportunities:

1. AFDP 3-0 Operations: Implement an evolving strategy to incorporate AI/Machine Learning algorithms in the overall planning and force projection processes. This would support the AI's training with real-world data and balance risk with the required operational tempo. Likewise, all downstream operational doctrines can employ AI to manage airmen, airspace, munitions and their combined employment.
2. AFDP 1-1 Mission Command and 3-60 Targeting: Two significant adjustments are required. If operating against an adversary enhanced by AI technologies, additional decision-making authority must be delegated to the lowest possible level in order to adjust quickly enough to hold effective engagements. Second, if blue forces are operating a friendly AI, decision making loops must specify where, when and how AI support will be used and where human-in-the-loop stopgaps will be employed. Similar to Mission Command, the Joint Targeting Cycle needs revision based upon the commander's intent in order to balance risk with tempo advantage.
3. AFDP 3-12 Cybersecurity and 3-85 EMS: Cybersecurity and Electronic Warfare will need to work together to protect the friendly use of AI while suppressing the effectiveness of the enemy. The four pillars established by the 2018 National Cyber Strategy would need to be expanded to specifically add offensive actions taken to deny enemy AI operations.

¹⁵ Department of the Air Force, Electromagnetic Spectrum Operations, Air Force Doctrine Publication 3-85 (Maxwell AFB, AL: Curtis E. LeMay Center for Doctrine Development and Education, December 14, 2023), 3