

FIGHTING TODAY'S PEER/NEAR-PEER FIGHT WITH TODAY'S TECHNOLOGY



Authors:

Major Heath Phillips (heath.phillips.2@us.af.mil)

Major Troy Taylor (troy.taylor.13@us.af.mil)

Major Robert Vincent (robert.vincent.1@us.af.mil)

All are currently students at Air Command and Staff College. Further biographical information will be supplied if needed for publication.

The conclusions and opinions expressed in this research paper are those of the authors and do not necessarily reflect the official policy or position of the U.S. Government, Department of Defense, or The Air University



1. Problem Statement

Centralized command and control (C2) is a luxury afforded to those with military superiority. Conflicts with near-peer adversaries will contest Jomini's interior lines of communication such that superiority may be local and temporary.¹ While centralized C2 presents less risk, the near-peer enemy forces a model of centralized command and decentralized control. The US Air Force's answer to the near-peer contest is to develop a mesh-networked C2 system called the Advanced Battle Management System (ABMS), bringing the "internet of things" to the battlefield.² Projected to cost hundreds of billions of dollars, ABMS is ambitious, creative, and nebulous.³ Before embarking on another massive acquisitions quest to build ABMS, akin to the F-22 and F-35 programs, the Air Force should first consider the following maxims. First, innovation does not necessarily require invention. In other words, consider the possibility that the hardware to meet the proposed capabilities and requirements that ABMS will address already exists and needs to be assembled, refined, and incrementally improved. Secondly, along this vein, the drive to counter improvised explosive devices led to tremendous computational advances in wide-area motion imagery (WAMI) for persistent intelligence, surveillance, and reconnaissance (ISR) packaged for use in low-cost attritable airframes. WAMI used for persistent ISR is the precursor to ABMS. Finally, the speed of future conflict may require the use of artificial intelligence within ABMS for automated targeting. The legal and ethical considerations of automated targeting must be considered prior to the acquisition of ABMS, before the Air Force opens a Pandora's Box of futuristic dystopia.

2. Defining the Urgency of Need

Today, China possesses the capability to render centralized aviation command and control (AC2) methods ineffective.⁴ To do so, they do not have to deny communication everywhere; they need only make processes too cumbersome to maintain, slowing down the Observe, Orient, Decide, Act (OODA) loop, which will force the U.S. to become reactive in nature.⁵ China will do this by being disciplined and quiet on our networks, manipulating data, and forcing the U.S. operator to lose faith in his or her C2 system, ultimately degrading its effectiveness.⁶ Next, or simultaneously, China will impede or degrade faster means of communications such as SATCOM and data services causing U.S. forces to execute their Primary, Alternate, Contingency, Emergency (PACE) plans.⁷ Logically, these plans tend to revert to less and less efficient communication mediums, i.e. high frequency (HF) forms of communication for long haul or beyond line of sight information exchange between C2 nodes.⁸ This drastically slows the feedback loop required for centralized AC2.

The U.S. military's near-term answer is to decentralize control to forward deployed C2 nodes and implement mission command that fosters commander's intent over detailed operations orders.⁹ Unfortunately, this assumes away the enemy's capability to deny access to even echeloned units of the Control and Reporting Center, Air Support Operations Center, or any service equivalent. The Chinese have focused extensively on Anti-Access Area Denial (A2AD) capabilities, building a defense in depth model that will stifle U.S. efforts to push operations forward.¹⁰ Defeating this is not impossible but will demand decentralization that enables faster than current human 'in the loop' targeting cycles. The question becomes, "does the military currently possess a capability that can achieve such an end?" Before answering this question, the reader must consider the specific requirements of such a capability.

3. Defining Requirements

Requirements begin in joint aviation doctrine. If one does not understand what is supposed to occur in planning, one cannot build an effective or ethically aligned AI-augmented, decentralized solution. In other words, processes the enemy will disrupt must still occur in some way inside of an AI-augmented, decentralized solution. To narrow the scope, this argument will only focus on the Joint Air Tasking Order (ATO) cycle from inception through execution and assessment, which is the Joint Force Air Component Commander's (JFACC's) means to lead theater targeting efforts.¹¹

The ATO Cycle is built on boards, bureaus, centers, cells, and working groups, or B2C2WG. All the following stages of the joint ATO Cycle are in some capacity mirrored in the air staff of individual service components--service cycles feeding the larger joint cycle. That said, ATO development begins with receipt of the Joint Force Commander's and JFACC's objectives, effects, and guidance.¹² This guidance takes the form of the Air Operations Directive (AOD), which ensures unity of effort among planners and the decentralized executors. Prior to the second stage, Target Development, the JFACC's staff convenes the first Joint Target Coordination Board (JTCCB), where all invested interests in targeting such as the Army, Navy, Marine Corps and special operations liaisons work through service specific needs. Target Development's output is the Joint Integrated Prioritized Target List which 'racks and stacks' the targeting effort. This prioritized list emphasizes targets conducive to operational level objectives, which intel analysts, lawyers, and planners have vetted throughout the Target Development and will continue to review into and through the third stage, Weaponizing and Allocation.¹³

The fourth stage builds and disseminates the ATO to operating units in theater. Information such as mission data, routing, controlling agency information, tanker plans and

fallout contingencies, target and weaponeering details all inform the battle as it is fought.¹⁴

However, a plan never survives first contact with the enemy. These four stages lay the framework and requirements of an AI-augmented, decentralized capability, but the fifth and sixth stages expose them explicitly. The system must possess three critical capabilities that come to light in the fifth and sixth stages of the ATO Cycle. Firstly, in the Execution Stage, it must understand and honor the aforementioned outputs of the various B2C2WG events. For example, it must comprehend the objectives laid out in the AOD, ascertain what actions or inactions will achieve the commander's end state, why certain targets matter more than others, and respect the ROE built into the targeting plan. Secondly, the solution system must possess the capability to dynamically communicate with aviation platforms to control airspace and integrate fires. And, it must demonstrate the ability to prioritize time sensitive targets and conduct dynamic collateral damage estimates prior to target prosecution. Lastly, the solution must have the capacity to begin the final Assessment stage of the ATO Cycle. This infers that the system can pull together battle damage assessments, bomb hit assessments, and understand and articulate what parts of the plan occurred and which did not. Logically, the system must be able to communicate these results back to centralized command nodes to inform planning efforts of future ATOs.¹⁵

4. Wide Area Motion Imagery Evolves into Command Nodes

Combating the threat of improvised explosive devices (IEDs) used by terrorist groups drove the development of WAMI for persistent ISR to become the most computationally powerful airborne asset to date.¹⁶ As the number of US casualties in Iraq and Afghanistan rose dramatically during the outset of Operations Enduring Freedom and Iraqi Freedom, the limitations of available persistent ISR became glaringly obvious. The Predator system with a single steerable camera suffered from the "soda straw" problem, meaning that the field of view

was so focused and limited that video analysts could not track multiple targets at once. A program named Constant Hawk filled this limitation by fusing a six-camera system together digitally into one enormous keystone image covering over 200 square kilometers of area with no gaps in coverage. Weather permitting, every event within the keystone was recorded for playback and analysis once the platform landed and the saved data was extracted. The program Angel Fire took this concept one step further by adding a radio frequency (RF) communications link, like a military grade Wi-Fi hub, to a nearby ground station so that ground forces could view the imagery feed in near-real time for immediate action. Angel Fire flew daily for almost two years over the city of Fallujah while streaming its imagery feed to the local Marine command post.¹⁷ This marked the technological beginning of the concept portrayed by J. R. R. Tolkien, the all-seeing Eye of Sauron for the battlefield.¹⁸

The computational horsepower required to process the sheer amount of data produced by the persistent cameras means that not only is there a watchful eye in the sky, but a brain as well. The successors to Angel Fire and Constant Hawk are Blue Devil and Gorgon Stare, both currently used heavily in Central Command (CENTCOM) operations.¹⁹ Both Blue Devil and Gorgon Stare can be packaged within UAVs, like the MQ-9 Reaper, and maintain a local data feed to ground forces along with global exfiltration through satellite networks. The automated analytical load to produce WAMI data is enormous. The Gorgon Stare package produces 65 trillion pixels of images in a 10-hour mission. Needless to say, this amount of imagery is taxing for human analysts to comb through and decipher. However, clever software development, also referred to as artificial intelligence (AI), presented a solution to this problem. WAMI now uses AI processing tools called "activity-based intelligence" to automatically assess adversarial behavior, originally used by football analysts to predict plays based on the line formation and

initial moments following the snap. Also, Gorgon Stare uses anomalous behavior detection routinely used by credit card companies to detect fraud to highlight unusual changes to a scene of interest.²⁰ The point being, AI assessments of combat zones by standalone airborne platforms via data fusion is already a reality. The path to automated control of nearby military assets is 90% complete. WAMI platforms now need to expand their networking capability to work in mesh tandem with additional WAMI platforms and establish ports and protocol links to other weapon systems, a step that is technically much less demanding than development up to this point. Once networked to regional systems and local ground troops, decentralized control is simply adding software. Awareness of this fact is crucial, lest the AF stumbles into an ABMS doppelganger without realizing what they created.

Similar to how self-driving autonomous cars are potentially safer than human-driven cars by a drastic margin, automated targeting using techniques like machine learning is potentially superior at avoiding civilian casualties and collateral damage.²¹ Once WAMI platforms network together in theatre and share computing resources, automated targeting is a capability readily loaded onto the WAMI processors. However, there are nuances to machine learning that must be explicitly acknowledged. War is both complicated and complex, so one cannot expect to preload definitive plans onto a decentralized automated C2 node. In other words, the art of war is too abstract to create a comprehensive physical model that a computer can chug through algebraically (yet!). Machine learning cleverly avoids this fact by completely parameterizing, and avoiding, the underlying character of war by training on the observable nature of the current conflict.²² All that machine learning requires are goals provided by the designer, digitized observables, and computational horsepower, both currently available with WAMI platforms. This is why machine learning marks a monumental shift in capability; the centralized command

will load their desired metrics and the decentralized control node will learn and adapt based on the ongoing conflict to dial in the commanded results.

Dependency upon training data presents a unique duality regarding precision and just war during conflict when using AI with machine learning. From the onset of the conflict, the decentralized control node may issue tactical plans that actually result in higher civilian casualties than would be coordinated by human military planners. This is because machine learning starts from an inaccurate initial condition. Consider figure 1 as a hypothetical example of civilian casualties with time comparing automated targeting with machine learning to conventional human in the loop targeting. At first, automated targeting may be shockingly poor, but with exponentially improving precision compared to modest incremental changes with conventional human targeting. The moral quandary lies in the space between the two curves in figure 1. Will, and should, the US accept higher civilian casualties initially knowing that far more civilians will be spared in the long run? Will the populous tolerate mistakes from a machine compared to those from a human? Who is morally and legally culpable when automated targeting gets it wrong? In summary, whether by intent or unguided evolution, the components of ABMS already exist and are assembling themselves together right now with WAMI and AI. A 20-year plan to build ABMS is like showing up to the race in a Ferrari, after the checkered flag.

5. Legal Considerations for Acquisition and Use of Autonomous Weapon Systems

An ongoing debate in the legal community revolves around some of the questions posed above. On one end of the debate, there are non-governmental organizations and policy groups advocating for a total ban of autonomous weapon systems (AWS) in warfare.²³ On the other hand, there are those who believe AWS may be employed now under current law, including the

Law of Armed Conflict (LOAC).²⁵ With sufficient checks in place, we argue that the U.S. should be able to develop and utilize fully AWS under existing international law and the LOAC.

Current DoD policy places a ban on all man out of the loop AWS. (DODD 3000.09, 4(a)-(c). This policy has been reinforced by former Defense Secretary Carter, who pledged that the DoD would never employ fully AWS with lethal capabilities.²⁶ As the U.S. has pivoted its focus toward near-peer competition, however, we have begun to realize the need for faster C2 constructs, especially in an environment in which communications are degraded.²⁷ Calls for bans and/or heavy restrictions on AWS are short sighted and ignore the fact that these weapon systems can comply with the LOAC such that they could be used lawfully today.

The DoD Law of War Manual provides guidance for legal reviews on new weapons and weapon systems to ensure such weapons comply with international law, specifically Article 36 of Additional Protocol 1 (API) of the Geneva Convention.²⁸ API states provides that any weapon system used in combat must not violate any principles of international law or treaties.²⁹ Given that AWS do not fall under the category of weapons prohibited by international law, the DoD conducts its legal review of new weapons or new technology applied to weapons, pursuant to the four guiding principles of the LOAC: namely, military necessity, distinction, proportionality, and humanity.³⁰

Military necessity is best described as “the principle that justifies the use of all measures needed to defeat the enemy as quickly and efficiently as possible that are not prohibited by the law of war.”³¹ Necessity is closely linked with distinction, and is the principle that “obliges parties to a conflict to distinguish principally between the armed forces and the civilian population, and between unprotected and protected objects.”³² The principle of proportionality prohibits any “attack which may be expected to cause incidental loss of civilian life, injury to

civilians, damage to civilian objects, or a combination thereof, which would be excessive in relation to the concrete and direct military advantage anticipated."³³ Finally, humanity “forbids the infliction of suffering, injury, or destruction unnecessary to accomplish a legitimate military purpose.”³⁴

Those calling for a ban or heavy restrictions on AWS argue the legal proposition that AWS cannot abide by these principles of the LOAC.³⁵ While others argue that the current international law construct provides a path for AWS usage in combat now, both sides must consider relevant questions. For example, can an AWS distinguish between a civilian and an enemy combatant (i.e., unlawful and lawful targets)? This distinction is incredibly difficult in certain scenarios, even for experienced military operators, especially in counterinsurgency operations in urban environments. Could an AWS be able to analyze new information in real time, such as performing complex decision-making tasks to determine whether damage caused destroying a given target would be excessive in relation to the direct military advantage gained by the attack? Again, this issue of proportionality is a difficult one to navigate, fraught with complex, subjective considerations difficult for even the most experienced commanders to make.

Notwithstanding the legal precautions taken with any weapon system, humans are prone to error. Humans also suffer from diminished decision-making capabilities in stressful, harmful environments when speed is necessary. Needless to say, humans make mistakes in targeting decisions frequently, and those mistakes are not always characterized as violations of the LOAC.³⁶ These issues bring up the issue of who should be held accountable should an AWS violate the LOAC or one of the ROEs, e.g. When humans violate principles of the LOAC, it is easy to hold someone accountable—either the individual making the mistake, or the commander making poor decisions. If an AWS violated the LOAC, though, should any individual be held

accountable? We would propose that commanders may be held accountable, pursuant to DoDD 3000.09, para. 4(b), which provides generally that commanders have a responsibility to operate and deploy AWS in a manner consistent with international law and with the capabilities the system possesses. This argument highlights another strength for employing AWS—namely, that AWS only do what they are programmed to do. To that end, programmers and developers should only program AWS with commands that are not capable of committing egregious violations of international law.³⁷

In conclusion, new weapons and new technologies have been met with resistance and calls for denunciation for hundreds of years.³⁸ The AWS envisioned in this paper encounters the same resistance, but can be shown to comply with international law and the LOAC to the extent that the weapons themselves are not particularly novel; rather, the AWS technology that allows for out of the loop decision-making is new. Despite DoD policy banning fully AWS, the benefits gained by employing such technology are sufficient to modify DoD policy in favor of systems that execute decisions much faster than any human, and have the ability to do so in areas of degraded communications capabilities. History has also shown that humans frequently make errors with weapon systems and that decision processes under stress and uncertainty are anything but reliable. While it may be beneficial in the near term to continue developing human in the loop AWS as we continue developing the technology, the U.S. should commit to developing and deploying fully AWS for use in a future near-peer conflict when we might be deprived of the luxury of reliable C2 networks.




Figure 1: Hypothesized relations of civilian casualties with time when using automated targeting with AI training during the onset of conflict compared to conventional human in the loop targeting with deliberate effort to reduce civilian casualties.

Notes

1. John Shy, “Jomini,” in *Makers of Modern Strategy*, ed. Peter Paret (Princeton, NJ: Princeton University Press, 1986), 170. Jomini is the first military strategist to suggest the importance of interior versus exterior lines of communication during war for command and control of dispersed forces. In short, disrupting interior lines of communication may be strategically decisive.

2. Sandra Erwin, “Air Force Paints a Digital Future Where Data From Satellites Play Central Role,” <https://spacenews.com/air-force-paints-a-digital-future-where-data-from-satellites-play-central-role/>, accessed 1 March 2020.

3. Ibid.

4. Missile Defense Advocacy, “China’s Anti-Access Area Denial,” <https://missiledefenseadvocacy.org/missile-threat-and-proliferation/todays-missile-threat/china-anti-access-area-denial-coming-soon/>, accessed 1 March 2020, Missile Defense Advocacy Alliance focuses their discussion on the China’s A2AD threat creating a defense in depth that puts maritime forces at risk if they choose to close the distance with China. The website paints Chinas threat picture by illustrating threat rings created by Chinese missile systems and maritime capabilities together. Clausewitzian theory on the strength of the defense emerges as a central theme in A2AD theory. Of note, the strength of Chinese defenses seems to be oriented solely toward the South China Sea.

5. Mark Pomerleau, “Breaking Down China’s Electronic Warfare Tactics,” (March 22, 2017); <https://www.c4isrnet.com/c2-comms/2017/03/22/breaking-down-chinas-electronic-warfare-tactics/>. Pomerleau discusses the emerging, and in many ways already mature EW capability resident in Chinese warfighters. Their plan, not unlike the United States is to use multiple effects across the cyber and space domains to create a holistic effect. An example is spoofing radars to confuse operators into losing faith in their systems while jamming communications in local areas.

6. Ibid.

7. Army Field Manual (FM) 6-02, *Signal Support to Operations*, September 2019, 1-7 – 1-8, <https://fas.org/irp/doddir/army/fm6-02.pdf>, This army publication walks through a typical communications degradation plan. Most importantly to this argument, FM 6-02 explains how as one works through the PACE plan, each step relies on less effective mediums. Various manuals in the Marine Corps and Air Force discuss the same planning considerations. The USMC DASC Handbook talks to communications planning considerations when fighting an enemy with known EW capabilities. All echo on another on this communication paradigm.

8. Ibid.

9. United States Marine Corps Concepts and Programs Manual, 2018, 9. The Marine Corps releases this manual to explain myriad concepts that it is seeking to employ or currently implementing across the operating forces. The Marine Corps has taken drastic steps to regain its sea roots as an expeditionary force from the sea—an intrinsic part of the Navy / Marine Corps team. As such, this manual addresses two pertinent issues throughout, Littoral Operations in a Contested Environment (LOCE) and Expeditionary Advance Basing Operations. Both rely on distributed C2 and decentralized nodes for mission accomplishment. Also reference the RAND study, “Distributed Operations in a Contested Environment: Implications for US Air Force Presentation.” The Joint Publication on distributed operations also indicates the movement

toward decentralization. A discussion on decentralized versus centralized C2 has been ongoing for years. Now, with near-peer competitors, arguments made by the likes of Lt Col Hinote in “Centralized Command Decentralized Execution: A Catchphrase in Crisis” for decentralization cannot be ignored.

10. Missile Defense Advocacy, “China’s Anti-Access Area Denial,” 1.

11. Joint Publication (JP) 3-30, *Joint Air Operations*, July 2019, III-22.

12. *Ibid.*, III-23.

13. *Ibid.*

14. *Ibid.*, III-24-25.

15. *Ibid.*, III-26.

16. Arthur Holland Michel, *Eyes in the Sky: The Secret Rise of GORGON STARE and How it will Watch us All*, (Boston, MA: Houghton Mifflin Harcourt, 2019), ch. 6. This work presents a thorough review of the origins of wide area motion imagery (WAMI) to combat the prolific use of improvised explosive devices (IEDs) used by terrorist and insurgent groups in Iraq and Afghanistan since 2002. It also summarizes the current state of WAMI systems and argues the point that the capability to surveil any given populous from above is shockingly advanced and has wide-ranging implications for security, privacy, and autonomy of computer systems.

17. *Ibid.*, ch. 1-2.

18. Jad Abumrad, Robert Krulwich, "Update: Eye in the Sky," *Radiolab* (podcast), 12 September 2016, <https://www.wnycstudios.org/podcasts/radiolab/podcasts>. This podcast discusses the implications of WAMI and how close the capabilities are to all-seeing status. The story focuses on how readily the applications shifted from finding IEDs to any number of security related issues in both the military and civilian sectors.

19. Michel, *Eyes in the Sky*, ch. 3.

20. *Ibid.*, ch. 7.

21. Lance Eliot, “Essential Stats for Justifying and Comparing Self-Driving Cars to Humans at the Wheel,” *Forbes*, 30 May 2019, <https://www.forbes.com/sites/lanceeliot/2019/05/30/essential-stats-for-justifying-and-comparing-self-driving-cars-to-humans-at-the-wheel/#5a1c09e746ed>. This article provides a summary of metrics available to assess human error while driving and shows the resulting statistics from these metrics. The point being that humans are surprisingly prone to error while operating vehicles, while the situation is unlikely to improve on its own over time.

22. Chris Meserole, *What is Machine Learning?* Brookings Report (Washington, DC: The Brookings Institution, 4 October 2018). This report provides a summary review of machine learning as a quantitative technique for the layperson.

24. For example, see Human Rights Watch (HRW), *Losing Humanity: The Case Against Killer Robots 2* (2012), which has stated that “Our research and analysis strongly conclude that fully autonomous weapons should be banned and that governments should urgently pursue that end.” Further, HRW has stated that “[HRW] finds that fully autonomous weapons would not only be unable to meet legal standards but would also undermine essential non-legal safeguards for civilians.” Report at <https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots>. See also the Campaign to Stop Killer Robots, which has declared that “Fully autonomous weapons would decide who lives and dies, without further human intervention, which crosses a moral threshold. <https://www.stopkillerrobots.org/learn/>.”

25. See, for example, Michael Press, “Of Robots and Rules: Autonomous Weapon Systems in the Law of Armed Conflict,” 48 *Georgetown Journal of International Law*, 1337 (Summer

2017), arguing generally that the LOAC provides sufficient legal justification for use of AWS in combat.

26. Sydney J. Freedberg and Colin Clark, “Killer Robots? ‘Never,’ Defense Secretary Carter Says,” *Breaking Defense* (blog), 15 September 2016, <https://breakingdefense.com/2016/09/killer-robots-never-says-defense-secretary-carter/>.

27. Press, “Of Robots and Rules,” 1342-43.

28. Office of General Counsel, Department of Defense, *Department of Defense Law of War Manual*, June 2015 (updated December 2016), para. 6.2, 337-338.

29. International Committee of the Red Cross, *Protocols Additional to the Geneva Conventions of 12 August 1949*, https://www.icrc.org/en/doc/assets/files/other/icrc_002_0321.pdf, Article 36, 30.

30. Maj Peter C. Combe II, “Autonomous Doctrine: Operationalizing the Law of Armed Conflict in the Employment of Lethal Autonomous Weapons Systems,” *51 St. Mary’s Law Journal*, 35 (2019).

31. *Law of War Manual*, para. 2.2.

32. *Law of War Manual*, para. 2.5.

33. *Additional Protocol I*, Art. 57(2)(iii).

34. *Law of War Manual*, para. 2.3.

35. Press, “Of Robots and Rules,” 1344.

36. See generally, Joel Hood, “The Equilibrium of Violence: Accountability in the Age of Autonomous Weapons Systems,” *11 BYU International Law and Management Review* 12, (Winter, 2015), arguing that humans are characterized as weapons systems as they operate various types of weapons, and that we make mistakes frequently. Such mistakes, he argues, do not make the weapon systems unlawful; rather the specific actions taken with those systems are what make them unlawful. To that end, the DoD should stringently word to ensure that programming on AWS is done in accordance with the LOAC and ROEs, and that commanders are trained to only use AWS for the purposes and missions they are designed, and have the capabilities to accomplish.

37. Press, “Of Robots and Rules,” 1360-1364.

38. *Law of War Manual*, para. 6.2.