



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2018-09-21

An Air Force Way of Swarm: Using Wargaming and Artificial Intelligence to Train Drones

Schuety, Clayton; Will, Lucas

University of Texas, National Security Network

<http://hdl.handle.net/10945/60592>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

An Air Force 'Way of Swarm': Using Wargaming and Artificial Intelligence to Train Drones

warontherocks.com/2018/09/an-air-force-way-of-swarm-using-wargaming-and-artificial-intelligence-to-train-drones/

September 21, 2018

Clayton Schuety and Lucas Will

September 21, 2018



For decades, swarms in nature, like locust swarms, have intrigued scientists and inspired them to uncover their fundamental laws. Swarms represent the ability to use the simple and many to accomplish the large and complex.

Now imagine being able to deploy a collection of drones on the scale and magnitude of a locust swarm. Indeed, the same laws that govern natural swarms could govern collective masses of cheap robotic systems. Today, efforts are underway to develop commercial swarms for crop pollination, surveillance, and high-resolution weather monitoring.

However, to date, developers largely centralize the management of their robotic swarms. Although this may work well in certain commercial applications, it may not necessarily work for military operations. Decentralized swarms are more resilient, given that centralized execution presents a single point of vulnerability that could be problematic in the contested environments of near-peer adversaries.

As two Air Force officers, we wonder, “is our service prepared to unleash the swarm?” We recently watched an official Air Force video of a hypothetical scenario, in which an F-35, teamed with a group of drones, launched on a mission to strike a strategic target deep inside the contested environment of a near-peer adversary. We wondered what would happen if dynamic threats emerged, forcing a departure from the pre-planned flight path, requiring the pilot to micromanage the decisions of each drone in an effort to salvage the strike. In such a case, the speed, scope, and complexity of the situation would overwhelm the pilot and the strike would fail. Although the F-35 had a capable team of drones, the key lesson was that a high number of drones alone does not equal a swarm.

For the Air Force, merely fielding a large group of drones in combat will not be decisive; the service must first figure out how to decentralize the drones’ execution and rebalance the workload of the human-autonomy teams. As a recent report by the Marine Corps Warfighting Lab shows, responsively controlling individual decisions of large numbers of drones is beyond the cognitive capabilities of a human. Instead, each drone must be able to execute the fundamental characteristic of the entire swarm, which is to independently coordinate its own decisions to produce behaviors to support a collective aim. Therefore, to effectively employ drones as a swarm, the human must delegate more freedom of action to the collective decision-making algorithms of their autonomous systems.

In this article, we extend the framework put forth by a recent *War on the Rocks* article describing “Athena,” a wargaming platform that tests artificial intelligence capabilities and captures and assesses the resulting data to improve military training. We argue that AI and wargames can be used to train not just human beings, but also drones. Specifically, by using narrow AI to play through millions of iterations of a mission-specific wargame, this framework uncovers the best rules for individual drone interaction that drive collective swarm behaviors supporting a specific mission. This “way of swarm” provides a method that is rapid in execution, highly flexible and adaptive to changing assumptions, and can focus on a specific mission while helping individual drones operate in decentralized fashion in service of a collective mission.

The Athena Approach

In their article on Athena, Benjamin Jensen, Scott Cuomo, and Chris Whyte contend that, instead of waiting for revolutions in general AI that will offer the flexibility and broad cognitive ability of humans, a focus on achievements in narrow AI, which is good at performing only specific tasks, offers a more immediate opportunity to solve specific military problems.

The authors write:

Wargames provide the optimal platform for exploring how to integrate AI with [commanders’] operational judgment. ... Our wargame, Athena, offers a way to build up a repository of data for future testing, enhance understanding of how AI can assist with training, red-teaming, and simulation, and highlight the limits of these capabilities as they interact with humans in uncertain environments.

We concur that the Athena framework of combining traditional wargaming with narrow AI can improve the quality of human training by reinforcing a commanders' operational and tactical decision-making process. Moreover, this framework can be applied to training drones. Narrow artificial intelligence, incorporated into wargaming platforms, presents a way to rapidly train quality decision-making algorithms to meet the Air Force's growing demand for swarming drones and loyal wingmen.

A training program for drone swarm algorithms that is rapid, flexible, and mission-specific will become paramount as the character of warfare changes. The confluence of commercially available technologies, like AI and drones, is progressing warfare into its next predicted evolution, where any force can now deploy a swarm into combat. Given the rise of long-term strategic competition between nations as well as their deliberate investment in AI and autonomy, figuring out how to not only leverage swarms, but also how to defend against them, is critical for a competitive strategy.

Insight from the Study of War

"Quantity has a quality of its own." – Attributed to Joseph Stalin

In a campaign for air dominance, a disadvantage in quantity can render the quality of the force irrelevant. In World War II, Russian forces blunted the German Air Force on the Eastern Front with sheer numbers. The same applied in the Pacific with the Japanese Air Force: While initially exceptionally well-trained, Japanese pilots quickly became too hard to replace in the quantities required to effectively compete. However, when two hypothetical forces are operating swarms of roughly similar sizes (as futurists anticipate), quality comes back into the picture. Simply possessing the same number of drones will no longer be decisive. Furthermore, in present-day swarms, each additional drone requires an increase in the workload of human controllers, eventually reaching a point of diminishing return for fielding a higher-quantity swarm. In these cases, success will come down to the quality – not the quantity – of the swarms employed.

Yet, what does quality mean in this context? Throughout history, it has often been not the hardware that was decisive in combat, but how well the human was trained to operate the technology (i.e., Flying Tigers, Sedan in 1940, Arab-Israeli Conflict, Falklands War). Studies on military training have found that the human was usually the limiting factor in the combat effectiveness of a weapon system, leading to a revolution in training programs across the services. During the Vietnam War, the Navy implemented Top Gun, a rigorous training program for their pilots that contributed to increasing the air-to-air kill ratio from two-to-one to twelve-to-one. After Vietnam, both the Air Force and Army learned from the Navy's success, and established large force exercises, like Red Flag and the National Training Center, to provide their operators with additional rigorous training cycles.

Besides just looking at how soldiers trained, the military began to focus on optimizing human decision-making processes. The Air Force trained aircrews to implement Colonel John Boyd's seminal model to speed the decision-making process by using a feedback loop of observing

the environment, orienting possible solutions, deciding based on the limited information, and acting to achieve a desired effect (the OODA loop).

“Never tell people how to do things. Tell them what to do and they will surprise you with their ingenuity.” – General George S. Patton, Jr.

The large force training exercises also helped the services enhance their doctrine of centralized control and decentralized execution. This concept became one of the fundamental tenets of the Air Force “way of war” as the scale, scope, and complexity of combat increased. The services now use the term mission command to embody the importance of decentralized execution in war.

Rigorous training, faster decision-making, and decentralized execution will continue to be critical when dealing with swarms of autonomous systems. Researchers contend that the decisive factor in the quality of autonomous systems will be the OODA loop-like algorithms inside the hardware that act either independently or in tandem with the human. But those algorithms need to be trained in order for the hardware to know what to do. In short, training humans will no longer be adequate; instead, training algorithms will become increasingly critical. Quality algorithms will enable more flexibility in the tactical environment and allow for decentralized execution at an unprecedented scale. Hence the need for a strategic framework that rapidly trains algorithms and builds trust and confidence between the human and the autonomous system.

Strategic Framework

One promising framework under development for training swarms of autonomous systems is DARPA’s Offensive Swarm-Enabled Tactics (OFFSET) program. OFFSET proposes to use a real-time game environment and a virtual reality interface to allow users to derive novel swarm tactics for autonomous systems through “crowd-sourcing” methods. By using mission-specific games to train, test, and employ swarming capabilities (rather than tailoring primitive swarming behaviors to a mission), the OFFSET framework displays significant promise.

However, we foresee three weaknesses with OFFSET. First, relying on crowd-sourcing efforts to experiment with the game may be difficult to sustain over time. The size of the “crowd” may not be sustainable as interest (and funding) in the project ebbs and flows. Second, as different hardware and environments become ready for testing, a crowd-sourcing method becomes cumbersome to rapidly repeat in training. Third, OFFSET is too slow. It emphasizes a real-time simulation environment for training swarms, where drones move via “point-and-click” through the battlespace. This approach reduces the speed and initiative of swarms in operations and makes it difficult to speed up repetitions to train drones through thousands of potential tactical scenarios in minutes or seconds.

However, combining OFFSET’s existing framework with narrow AI can produce a novel framework that is rapid, flexible, and adaptive to a variety of missions. In 2017, Google’s Deep Mind surprised many when it applied a similar “self-play” training framework (i.e. reinforced machine learning) to generate an algorithm that mastered the Chinese game of Go. Google’s

narrow AI was able to accumulate thousands of years of human knowledge during a period of just a few days, which was only possible through faster-than-real-time simulations. Combining Google's breakthrough in narrow AI with mission-specific wargames provides a powerful way to train and fight with swarms of drones.

For example, narrow AI could play a wargame that simulates high-level decisions, actions, interactions, and resulting behaviors for a swarm of aerial drones tasked to defend a base. Instead of solving for a centralized solution to manage the placement of the entire defending swarm, narrow AI iteratively discovers the best rules for individual interactions that, combined, generate a collective swarm behavior that minimizes base damage from an attacking force. Ultimately, the resulting "AI-trained" local interaction rules get loaded into each drone in a real-world swarm, ready to perform the specific base defense mission. The optimized local drone interaction rules enable self-organization and decentralization, which reduces human oversight to perform the specific mission.

We argue that this proposed Air Force "way of swarm" addresses the weakness of current drone training frameworks. First, it is rapid in execution and insulated from the instability of crowd-sourcing. The workload demand of a human crowd is replaced by the persistent availability of AI algorithms and cloud computing. Second, both the narrow AI and wargame would be highly flexible and adaptive to changing assumptions of the game, such as new hardware or environmental conditions. For instance, if a new sensor became available allowing the drone to better detect the adversary, the framework could rapidly re-run to again master the modified game. Third, this framework continues to apply the principle of using missions to determine capabilities rather than the other way around, but it does so by allowing the narrow AI to solve local drone interaction rules that preserve the principles of decentralized execution.

Anticipated Challenges

Combining high-fidelity wargames and cutting-edge narrow AI requires integrating private sector knowledge and military subject matter experts. Unfortunately, friction around this type of partnership has already occurred with nearly 4,000 Google employees demanding an end to their company's support to the Defense Department over Project Maven. That project aimed to use narrow AI to reduce the human workload and minimize non-combatant casualties, by helping military analysts better process, exploit, and disseminate the vast amount of collected intelligence, surveillance, and reconnaissance data.

Additionally, when using narrow AI there is a risk of the "black box" phenomenon. This occurs when there is no simple explanation for the decision an algorithm makes due to the complexity of the machine learning technique used. Ultimately, the military must strike a balance in how much risk it assumes. The trade-off is between low human bias and more decision-making flexibility, or high human bias that offers greater insight into why an algorithm made a particular decision, which can increase trust and confidence.

"No war is over until the enemy says it's over. We may think it over, we may declare it over, but in fact, the enemy gets a vote." – Secretary of Defense Jim Mattis

Some have even argued that the swarm may be “dead on arrival.” First, these critics point out, drones may not be as cheap as those advocating for the technology predict. But while we acknowledge that weaponizing drones may be cost-prohibitive at this point, commercially available systems like the DJI Mavic 2 (widely used by U.S. soldiers until a recent ban) are cheap enough (approximately \$1,000 each) to achieve a sufficient quantity to leverage the benefits of a swarm in non-kinetic missions, given the right decision-making algorithms. Second, swarm skeptics also claim that the same information technology that will enable drone swarms will also enable an adversary to deploy a cheap and effective air defense to defeat them. Although there is evidence to support this claim, history reminds us that even the best air defenses aren’t always perfect. For example, Stanley Baldwin’s proclamation in 1932 that “the bomber will always get through” an air defense proved to be true during World War II *if* the bomber’s force was willing to sustain high casualties. Swarms of cheap autonomous systems, by nature, reduce the risk of sustaining high casualties. Therefore, with high enough numbers, the swarm will always get through, despite how technologically advanced the air defense is.

Finally, some opponents of computer-based wargames contend that these models will produce an gap between theory and reality because they cannot capture human motivations like desire, commitment, passion, or will. Although these critiques hold merit, it also depends on the intended use of the wargame; the more specific (less generalizable) the wargame, the more the model is representative. Moreover, validating wargames with additional methods, such as field testing, can strengthen their utility for prediction and reduce the gap between simulations and reality. In the case of field testing swarms, the data could feed back into the AI and wargaming framework to generate more refined behaviors in subsequent operations.

A Way of Swarm

Many have recognized the rising cost of military hardware and the corresponding decrease in the quantity the services can field. The military has sought to reverse this trend by investing in large quantities of cheap systems (like drone swarms), yet it has overlooked a crucial element: how to train the swarm. To leverage the full potential of a swarm, the Air Force must build trust within the human-autonomy team and train to decentralize swarm execution.

We argue that this Air Force “way of swarm” holds promise for training a mastered array of tactics for its upcoming swarms of autonomous systems. Using narrow AI to play mission-specific wargames provides a rapid, flexible, and adaptive way to discover the best local interaction rules leading to useful collective swarm behaviors; this enhances human-autonomy teams by taking the workload off the human. Without delegating more freedom of action to the autonomous systems of a swarm, the speed, scope, and complexity of the future operating environment will outpace human-autonomy teams. Just as efforts to train quality aircrews were strategically decisive in the past, in the future, training the best quality algorithms for the Air Force’s drone swarms and teaming them with airmen will help the U.S. military maintain a competitive military advantage.

Major Clayton Schuety is a career gunship pilot with Air Force Special Operations Command. He is currently studying Irregular Warfare in the Defense Analysis Department at the Naval Postgraduate School in Monterey, California, working towards his second master's degree. He is a graduate of the United States Air Force Weapons School and has flown the AC-130H Spectre and the AC-130W Stinger II gunships while deployed in support of Operation Enduring Freedom, Operation Inherent Resolve, and multiple other contingency operations.

Major Lucas Will is a senior pilot flying remotely-piloted aircraft with Air Force Air Combat Command. He is currently studying Irregular Warfare in the Defense Analysis Department at the Naval Postgraduate School in Monterey, California, working towards his second master's degree. He has flown the MQ-1B Predator remotely-piloted aircraft and the HH-60G combat search and rescue helicopter in direct support of Operation Enduring Freedom, Operation Inherent Resolve, and multiple other contingency operations and humanitarian relief efforts.

The views expressed here are theirs alone and do not necessarily reflect those of the U.S. Air Force.

Image: [U.S. Marines/Sgt. Lucas Hopkins](#)

[Commentary](#)