



The upside and downside of swarming drones

Irving Lachow

ABSTRACT

The US and Chinese militaries are starting to test swarming drones – distributed collaborative systems made up of many small, cheap, unmanned aircraft. This new subset of independently operating or "autonomous" weapons is giving rise to new strategic, ethical, and legal questions. Swarming drones could offer real advantages, including reducing the loss of both human life and expensive equipment in battle. But they also come with potential dangers. There is already great international concern about deploying weapons without "meaningful human control." Proliferation is another danger, and a problem that could be particularly acute in the case of swarming drones. The risks posed by swarming drones should be considered sooner rather than later, before their destructive potential reaches maturity.

KEYWORDS

Distributed collaborative systems; drones; swarming drones; LAWS; lethal autonomous weapon systems; UAS; unmanned aircraft systems

Even as debate surrounding the use of autonomous weapons swirls, a nascent subset of the technology is giving rise to new strategic, ethical, and legal questions. Swarming drones, also known as distributed collaborative systems, are flocks of small unmanned aerial vehicles that can move and act as a group with only limited human intervention. Last October, the US Defense Department demonstrated what it said was one of the world's largest micro-drone swarms. It launched a flock of 103 Perdix drones into the sky above California, where they flew in formation and demonstrated collective decision-making without human help (Department of Defense 2017).

Swarming drones – which have not yet been used in warfare but are being tested and developed by other militaries – come with a whole range of implications about what armed forces can and should do with the technologies at their disposal. They could have real advantages, including reducing the loss of both human life and expensive equipment in battle. But they also come with potential dangers. The risks posed by swarming drones should be considered sooner rather than later, before their destructive potential reaches maturity.

Just what is a killer robot?

There is already plenty of international concern over unmanned weapon systems in general. In April 2013, Human Rights Watched launched its campaign to Stop Killer Robots (Human Rights Watch n.d.). In July 2015, some of the world's top artificial intelligence (AI) and robotics researchers released an open letter calling for "a ban on offensive autonomous weapons beyond meaningful human control" (Future of Life Institute 2016). While military research on autonomous and distributed collaborative systems continues unabated, there is debate within the US armed services about what level of human control is appropriate (Magnuson 2016).

It's impossible to discuss the legal and ethical implications of autonomous weapons, swarming or otherwise, without defining them, and doing so is harder than it might at first appear. In fact, there is little agreement on how to describe and categorize the key technologies that are the topics of this paper. For example, the term "unmanned aircraft system" or UAS, often used by the Federal Aviation Administration, encompasses both human-controlled drones and fully autonomous swarms because it focuses on whether the systems in question are manned, not whether they have the capacity to be fully independent, autonomous devices. The term says nothing about whether an unmanned aircraft is performing lethal or nonlethal missions, or taking an offensive, defensive, or supporting role in military operations. (One of the most common uses of UASs in both military and civilian capacities is to simply provide surveillance and reconnaissance.)

In contrast, the term used by the United Nations in many of its current studies, lethal autonomous

weapons systems (LAWS), excludes human-piloted drones but includes "guns" that can be operated without human intervention, like the US Navy's Phalanx and Israel's Iron Dome (Chansoria 2015). The focus on lethality limits the UN discussion to systems that either attack or defend. Intelligence, surveillance, and reconnaissance roles are excluded when framing the problem this way, which could omit swarming drones - such as the Defense Advanced Research Projects Agency's Gremlins project - that do intelligence-gathering (Adams 2016). This may be an important oversight, because some scholars are concerned about the impact autonomous systems could have on international human rights through their ability to deploy surveillance capabilities that undermine privacy (Roff 2015).

Moreover, there exists no agreement on how to define "autonomous," because various degrees of autonomy are possible. As Nicolas Marsh, a researcher at the Peace Research Institute Oslo, wrote, "setting the threshold of autonomy is going to involve significant debate because machine decision-making exists on a continuum" (Marsh 2014). For example, some systems are autonomous in some respects but not others; that is, a system may be able to autonomously detect a target, but not fire upon it without human intervention. Similarly, a weapon system that is defending an area may be granted permission to fire only on an incoming target that exhibits certain characteristics, like traits exclusive to missiles. Finally, there are systems that are launched by humans but then select and destroy their targets autonomously. In fact, such "fire and forget" systems have been used for many years, with dozens deployed by multiple countries (Marsh 2014). While the United States has fielded lethal weapons systems that are capable of operating autonomously, such systems operate with a human in the loop in accordance with US policy (Schmitt 2013; Department of Defense 2012).

Dawn of the mechanical swarm

Swarming behavior is most often seen in nature, such as when a school of fish or a flock of birds rapidly change direction in unison, in what looks like a series of tightly choreographed maneuvers. Swarming systems typically consist of individual agents (such as ants, birds, cells, or unmanned aircraft) that interact with one another and their environment. The agents follow simple rules, but the collective interactions between the agents can lead to quite complicated and sophisticated collective behaviors, including emergent intelligence. For example, a swarm may stay in formation while changing direction multiple times.

John Arquilla and David Ronfeldt, who were affiliated with the RAND Corporation at the time, wrote the first seminal work on the idea of swarming in the military arena in 2000. They defined it as "systematic pulsing of force and/or fire by dispersed, interneted units, so as to strike the adversary from all directions simultaneously" (Arguilla and Ronfeldt 2000, 8). Despite the value of their work, their definition is too limiting for the many purposes to which swarming can be applied in warfare. A more useful explanation is provided by Paul Scharre of the Center for a New American Security, who defined swarming as "large numbers of dispersed individuals or small groups coordinating together and fighting as a coherent whole" (Scharre 2014, 26). Swarms can exhibit emergent intelligence by following simple rules that guide the behavior of individual members of the swarm. When a swarm's agents follow these rules, complex and unified behaviors can arise. For this to occur, agents in the swarm must be homogeneous: They must have the same physical characteristics, the same programming, and the same sensors. Sensors are important because the rules used to guide swarm behavior are often based on environmental factors outside of the swarm. Finally, the agents in an autonomous swarm must be able to communicate with each other.

In nature, the size of a swarm can vary by many orders of magnitude - swarm intelligence can be displayed by a flock of a few dozen birds or by millions of microbes. The current state of the art in military applications is limited to swarms ranging in size from tens to hundreds of unmanned aircraft. This range seems to be a sweet spot for creating swarms, because the drones can be quite capable while still being much cheaper than their manned equivalents. The US Navy's LOCUST program - short for Low-Cost UAV Swarming Technology - follows this approach (Hambling 2016a). LOCUST's goal is to have 30 drones flying together like a flock of birds. A human pilot controls the behavior of the swarm, but does not fly the individual drones; they maintain their formation automatically. At around \$500,000 for a 30-drone swarm, the cost of LOCUST is less than half the price of the million-dollar Harpoon antiship missile it seeks to replace.

Looking further into the future, larger swarms are certainly possible. According to Zheng Dengzhou, an executive at the China Electronics Technology Group Corporation, "currently one man can control multiple drones in the system. In the future, one man will be

able to control hundreds and thousands of drones" (Hambling 2016b). Scharre believes that the military could create swarms of millions, perhaps even billions, of tiny, ultracheap 3-D-printed mini-drones. To prove his point, he cites a mini-drone created by researchers at the Harvard Microrobotics Lab; called "Mobee," it is manufactured using a 3-D printer that produces sheets of 2-D forms that can then be folded to create 3-D, bug-like vehicles. Commercialization of such minidrones is probably 5-10 years away (Ravindran 2014), but the development of such systems now seems inevitable.

Military uses

Swarming drones can be used in three ways by military forces: to attack, defend, and provide support functions such as intelligence, surveillance, and reconnaissance (Scharre 2014). Swarming is advantageous for offensive missions because it can overwhelm enemy defenses with a large number of potential targets. In a swarming attack the drones are dispersed, which makes it difficult and expensive for the adversary to defend itself. If 10 drones attack a target simultaneously and 7 are shot down, 3 will still be able to complete their mission. Because individual drones in a swarm do not need to survive at a high rate, they may be dramatically cheaper than standalone weapon systems, which are often extremely sophisticated and expensive. It is possible that even a large swarm of several dozen unmanned aircraft may be both more effective and less expensive than a single manned or unmanned aircraft.

Though overwhelming attention is paid to drones' potential offensive uses, they may be just as useful for defense, especially if one's adversary is using swarming tactics. The US Navy has run simulations to determine the ability of ships to protect themselves from swarming drone attacks, and the results are sobering: If eight drones attack a ship, then on average, three penetrate its defenses - including such advanced automated antidrone weapons as the Phalanx (Hambling 2016a). The outcomes become much worse when the number of attacking drones increases. Not surprisingly, the Navy is currently doing research on using defensive swarms to halt attackers (Hambling 2016a). In fact, a swarm of drones may be an effective way to protect oneself from a swarm of drones: drone against drone, if you will.

Swarms of drones can also be used for defensive purposes by creating large numbers of decoys that sow confusion or actively disrupt an attacking force. Scharre describes how miniature air-launched decoys can be used to fool enemy radars. He also notes that large numbers of drones could swarm over an enemy's airfield to prevent aircraft from taking off. A similar tactic could be used to protect a piece of territory from overflights by enemy helicopters or airplanes, though the difficulty of such a mission would increase with the size of the area that needed protecting.

Command and control

One of the fears associated with lethal autonomous weapon systems is that they might select and attack targets without meaningful human control (Human Rights Watch 2016). Swarms can further exacerbate such concerns because they may be able to attack in ways that are "intelligent," even though the elements of the swarm are following simple rules. What if a swarm "learns" in ways that were not anticipated, or develops an unforeseen level of sophistication? When a system is reacting in real time to a dynamically changing environment, and basing its decisions on simple sets of rules, it is possible that unanticipated behaviors will naturally arise. One must also remember that adversaries may actively try to undermine a swarm's effectiveness by fooling the drones, hacking their sensors or jamming their communications links (Russon 2015).

While these risks are real, though, the likelihood of man-made swarms running amok is not as high as many fear, for several reasons. First, when swarms are used in a military context, they are programmed to accomplish a specific mission or task as part of a broader plan - their range of possible behaviors can be bounded. The US military makes every effort to maintain command and control over swarms of drones, just as it does for other weapon systems (Scharre 2014, 35). The precise level of control, and of swarm autonomy, will vary by mission; however, military doctrine requires a human controller to ensure that a swarm follows proper rules of engagement and principles of humanitarian law (Department of Defense 2012). According to Michael Schmitt, a noted scholar on international humanitarian law and a professor at the US Naval War College, "As presently envisaged, autonomous weapon systems will only attack targets meeting predetermined criteria and will function within an area of operations set by human operators" (Schmitt 2013, 6). To address the risks of unauthorized actions due to jamming or hacking that may affect command and control, drones can be programmed with default rules of behavior if they lose contact with each other or their human overseer.

Legal and ethical considerations

The heart of the debate surrounding lethal autonomous weapon systems is largely about the legal and ethical considerations associated with the use of systems that may be outside of human control. Legal discussions focus primarily on the ability of automated systems to implement the key principles of international humanitarian law: distinction, necessity, and proportionality (Schmitt 2013). There are two key questions. First, can the autonomous weapon reliably distinguish, under difficult battlefield circumstances, between combatants and civilians - and only target the former? Second, can the autonomous weapon apply only the amount and kind of force necessary to defeat the enemy, without going further and causing excessive loss of civilian life?

Currently, military operations involving humancontrolled systems follow international humanitarian law. This is a fundamental tenet of the US policy and that of its key allies, and the US Government intends that autonomous systems follow the same principles (Department of Defense 2012). The main question up for debate is whether fully autonomous systems can do so. The answer depends on what the term "fully autonomous" means - what functions are automated? - and whether such a system is still under some level of human control. It also depends on the intentions of the party deploying the autonomous weapon system. If a government wants its autonomous weapons to follow international humanitarian law, it can take steps to make sure they do so. Finally, the answer depends on the technical capabilities of the system and whether it can exhibit desirable behaviors when confronted with novel and complex situations.

A second legal framework that may apply to the debate surrounding autonomous weapon systems is international human rights law, which enters the picture if these systems are used for surveillance that could compromise privacy.

The primary ethical debates around autonomous weapons focus on two main issues: First, is it more or less moral to employ autonomous rather than humancontrolled weapons? To answer this question, one must weigh the potential benefits that come from using unmanned systems, such as better decision-making and reduced loss of life for the party operating them, against the potential harm that could be caused by them, such as removal of human emotion from decision-making and an increased risk of collateral damage. Second, is the use of autonomous systems in warfare itself an immoral act that should be precluded? UN Human Rights Council Special Rapporteur Christof Heyns has argued that it is an affront to human dignity to allow machines to make life and death decisions (UNIDIR 2015, 7).

A final ethical issue to consider is the potential future impact of autonomous weapon systems due to proliferation - a problem that could be particularly acute in the case of swarming drones. Both scholars and military officials have expressed concern that if the United States develops autonomous weapons today, it may lead to their proliferation and use by adversaries with fewer moral concerns. "Different cultures may have very different views on what is perceived as ethical and moral ... It is possible therefore that we may face an adversary in the future whose freedom of action, with regard to autonomy, is considerably greater than our own," wrote Group Captain Clive Blount of Great Britain's Royal Air Force (Blount n.d., 38). Autonomous weapons could, of course, be extremely useful to terrorists and repressive governments. Given this risk, some believe it is more ethical to forego their potential benefits to minimize the potential harm that could come from rogue actors using the same technologies (Human Rights Watch 2016). The counterargument is that bad actors might be able to deploy these types of systems anyway. If they do, then it may be unwise for countries like the United States to forego having the same technologies for defensive purposes.

Big decisions

The ultimate question that scholars, scientists, and policy-makers are grappling with is whether lethal autonomous weapon systems should be banned, regulated, or allowed without restraint. Unsurprisingly, these are not easy issues to analyze. Difficult questions await: If there is to be a ban, when should it happen preemptively or after these weapons are better known and understood? To what, exactly, should the ban be applied? Will it prevent the deployment of defensive automated systems as well as offensive ones? Will countries agree to such a ban? If some don't but some do, will the latter group be at a military disadvantage? How will a ban be enforced?

If the path forward is regulation rather than a ban, similar questions arise. What exactly will be regulated? Who will decide what regulations should be applied, and who will enforce them? What will the consequences for violators be? Attempts to answer these questions involve complex debates about operational effectiveness, technical maturity, legal rights, and ethical perspectives.

Swarming drones are particularly difficult to address because of the offensive-defensive dynamic they present. Early experiments indicate that the only effective protection against attacking swarms may be defensive swarms (Hambling 2016a). If this turns out to be true, then advanced militaries will be incentivized to develop swarms. Here is how the dynamic could play out: If defensive swarms are deployed, then attackers may feel the need to create offensive swarms that are larger and more capable than the defensive swarms they will face. If, meanwhile, offensive swarms are deployed, then defenders will need to deploy defensive swarms because that is the only way they will be able to defend themselves. It is easy to see how swarm development and deployment could lead to an arms race focused on these technologies.

At present, two notable efforts are underway to help divine the way forward. In October 2016, the White House issued a report that addresses the future of AI. Looking at the use of AI in weapon systems, the report affirmed that "all weapon systems, autonomous or otherwise, must adhere to international humanitarian law, including the principles of distinction and proportionality" (Executive Office of the President 2016, 37-38). However, the report also acknowledged that due to the rapid pace of technological development, the United States must develop a government-wide policy on the use of autonomous weapons that is "consistent with shared human values, national security interests, and international and diplomatic obligations" (Executive Office of the President 2016, 38).

Internationally, the United Nation's Convention on Certain Conventional Weapons (CCW) is examining the issue of LAWSs. The CCW's purpose is to "ban or restrict the use of specific types of weapons that are considered to cause unnecessary or unjustifiable suffering to combatants or to affect civilians indiscriminately" (United Nations Office at Geneva n.d.). The CCW has created a Group of Government Experts to study the issue and make recommendations (Sukumar 2016). While the group's recommendations will not be legally binding, they can serve as a set of norms that influence international behavior. This type of expertgroup process normally takes a few years, which means that a decision about the fate of autonomous weapons systems, including swarming drones, is not imminent. In the meantime, the United States has the opportunity to develop a government-wide policy and continue exploring the benefits and risks of these systems. Swarming drones bring many benefits, but they also carry risks that need to be addressed for the good of all.

Disclosure statement

The author's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by the author.

Notes on contributor

Irving Lachow is the portfolio manager for international cyber programs at the MITRE Corporation. He is also a visiting fellow at the Hoover Institution and an affiliate at the Center for International Security and Cooperation at Stanford University. He received his PhD in engineering and public policy from Carnegie Mellon University.

References

Future of Life Institute. 2016. "Autonomous Weapons: An Open Letter from AI and Robotics Researchers." July 28. Accessed January 7, 2017. http://futureoflife.org/openletter-autonomous-weapons/

Adams, E. 2016. "DARPA's Developing Tiny Drones that Swarm to and from Motherships." April 13. Accessed January 7, 2017. https://www.wired.com/2016/04/darpasdeveloping-tiny-drones-swarm-motherships/

Arquilla, J., and D. Ronfeldt. 2000. Swarming & the Future of Military Conflict. DB311. Santa Monica, CA: RAND.

Blount, Group Captain Clive. n.d. "War at a Distance? -Some Thoughts for Airpower Practioners." Air Power Review 31-39.

Chansoria, M. 2015. "Autonomous Weapons: Tightrope Balance." Bulletin of the Atomic Scientists, November

Department of Defense. 2012. Autonomy in Weapon Systems. Directive, Washington, DC: Department of Defense.

Department of Defense. 2017. Department of Defense Announces Successful Micro-Drone Demonstration. Press release number NR-008-17. January 9. https://www. defense.gov/News/News-Releases/News-Release-View /Article/1044811/department-of-defense-announcessuccessful-micro-drone-demonstration

Executive Office of the President. 2016. Preparing for the Future of Artficial Intelligence. Washington, DC: White House: Committee on Technology, National Scient and Technology Council.

Hambling, D. 2016a. "U.S. Navy Plans to Fly First Drone Swarm This Summer." Defense Tech, January 4. http:// www.defensetech.org/2016/01/04/u-s-navy-plans-to-flyfirst-drone-swarm-this-summer/

Hambling, D. 2016b. "If Drone Swarms are the Future, China May Be Winning." Popular Mechanics, December 23. http://www.popularmechanics.com/military/research/ a24494/chinese-drones-swarms/

Human Rights Watch. 2016. "UN: Key Action on 'Killer Robots." December 16. https://www.hrw.org/print/ 297850



Human Rights Watch. n.d. "Killer Robots." Accessed January 7, 2017. https://www.hrw.org/topic/arms/killer-robots

Magnuson, S. 2016. "Military Beefs Up Research Into Swarming Drones." March 21. Accessed January 7, 2017. http://www. nationaldefensemagazine.org/archive/2016/March/Pages/ MilitaryBeefsUpResearchIntoSwarmingDrones.aspx

Marsh, N. 2014. Defining the Scope of Autonomy. PRIO Policy Brief. Oslo: Peace Research Institute Oslo.

Ravindran, S. 2014. "Insect-Inspired Vision Helps These Tiny Robots Fly." June 17. Accessed December 26, 2017. http://motherboard.vice.com/read/insect-inspired-visionhelps-these-tiny-robots-fly

Roff, H. 2015. "Banning and Regulating Autonomous Weapons." Bulletin of the Atomic Scientists, November 24.

Russon, M.-A. 2015. "Wondering How to Hack a Military Drone? It's All on Google." International Business Times, May 8. http://www.ibtimes.co.uk/wondering-how-hackmilitary-drone-its-all-google-1500326

Scharre, P. 2014. Robotics on the Battlefield Part II: The Coming Swarm. Washington, DC: Center for a New American Security.

Schmitt, M. 2013. "Autonomous Weapon Systems and International Humanitarian Law: A Reply to Critics." Harvard National Security Journal, February 5.

Sukumar, A. M. 2016. "India to Chair UN Group on 'Killer Robots', Open New Page on Arms Control Diplomacy." The Wire, December 19. Accessed January 2, 2017. https://thewire.in/87890/india-chair-un-groupkiller-robots-open-new-page-arms-control-diplomacy/

UNIDIR. 2015. The Weaponization of Increasingly Autonomous Technologies: Considering Ethics and Social Values. New York, NY: United Nations: UNIDIR Resources. No. 3.

United Nations Office at Geneva. n.d. "The Convention on Certain Conventional Weapons." Accessed January 1, 2017. http://www.unog.ch/80256EE600585943/(httpPages)/ 4F0DEF093B4860B4C1257180004B1B30?OpenDocument