Pros and Cons of Autonomous Weapons Systems

Amitai Etzioni, PhD

Oren Etzioni, PhD

Autonomous weapons systems and military robots are progressing from science fiction movies to designers' drawing boards, to engineering laboratories, and to the battlefield. These machines have prompted a debate among military planners, roboticists, and ethicists about the development and deployment of weapons that can perform increasingly advanced functions, including targeting and application of force, with little or no human oversight.

Some military experts hold that autonomous weapons systems not only confer significant strategic and tactical advantages in the battleground but also that they are preferable on moral grounds to the use of human combatants. In contrast, critics hold that these weapons should be curbed, if not banned altogether, for a variety of moral and legal reasons. This article first reviews arguments by those who favor autonomous weapons systems and then by those who oppose them. Next, it discusses challenges to limiting and defining autonomous weapons. Finally, it closes with a policy recommendation.

Arguments in Support of Autonomous Weapons Systems

Support for autonomous weapons systems falls into two general categories. Some members of the defense community advocate autonomous weapons because of military advantages. Other supporters emphasize moral justifications for using them.

Military advantages. Those who call for further development and deployment of autonomous weapons systems generally point to several military advantages. First, autonomous weapons systems act as a force multiplier. That is, fewer warfighters are needed for a given mission, and the efficacy of each warfighter is greater. Next, advocates credit autonomous weapons systems with expanding the battlefield, allowing combat to reach into areas that were previously inaccessible. Finally, autonomous weapons systems can reduce casualties by removing human warfighters from dangerous missions.[1]

The Department of Defense's *Unmanned Systems Roadmap: 2007-2032* provides additional reasons for pursuing autonomous weapons systems. These include that robots are better suited than humans for "dull, dirty, or dangerous' missions."[2] An example of a dull mission is long-duration sorties. An example of a dirty mission is one that exposes humans to potentially harmful radiological material. An example of a dangerous mission is explosive ordnance disposal. Maj. Jeffrey S. Thurnher, U.S. Army, adds, "[lethal autonomous robots] have the unique potential to operate at a tempo faster than humans can possibly achieve and to lethally strike even when communications links have been severed."[3]

In addition, the long-term savings that could be achieved through fielding an army of military robots have been highlighted. In a 2013 article published in *The Fiscal Times*, David Francis cites Department of Defense figures showing that "each soldier in Afghanistan costs the Pentagon roughly $850,000 per year."[4] Some estimate the cost per year to be even higher. Conversely, according to Francis, "the TALON robot—a small rover that can be outfitted with weapons, costs $230,000."[5]

As autonomous weapons systems move from concept to reality, military planners, roboticists, and ethicists debate the advantages, disadvantages, and morality of their use in current and future operating environments. (Image by Peggy Frierson)

According to *Defense News,* Gen. Robert Cone, former commander of the U.S. Army Training and Doctrine Command, suggested at the 2014 Army Aviation Symposium that by relying more on "support robots," the Army eventually could reduce the size of a brigade from four thousand to three thousand soldiers without a concomitant reduction in effectiveness.[6]

Air Force Maj. Jason S. DeSon, writing in the *Air Force Law Review*, notes the potential advantages of autonomous aerial weapons systems.[7] According to DeSon, the physical strain of high-G maneuvers and the intense mental concentration and situational awareness required of fighter pilots make them very prone to fatigue and exhaustion; robot pilots, on the other hand would not be subject to these physiological and mental constraints. Moreover, fully autonomous planes could be programmed to take genuinely random

and unpredictable action that could confuse an opponent. More striking still, Air Force Capt. Michael Byrnes predicts that a single unmanned aerial vehicle with machine-controlled maneuvering and accuracy could, "with a few hundred rounds of ammunition and sufficient fuel reserves," take out an entire fleet of aircraft, presumably one with human pilots.[8]

In 2012, a report by the Defense Science Board, in support of the Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, identified "six key areas in which advances in autonomy would

have significant benefit to [an] unmanned system: perception, planning, learning, human-robot interaction, natural language understanding, and multiagent coordination."[9] *Perception*, or perceptual processing, refers to sensors and sensing. Sensors include hardware, and sensing includes software.[10]

Next, according to the Defense Science Board, *planning* refers to "computing a sequence or partial order of actions that … [achieve] a desired state."[11] The process relies on effective processes and "algorithms needed to make decisions about action (provide autonomy) in situations in which humans are not in the environment (e.g., space, the ocean)."[12] Then, *learning* refers to how machines can collect and process large amounts of data into knowledge. The report asserts that research has shown machines process data into knowledge more effectively than people do.[13] It gives the example of machine learning for autonomous navigation in land vehicles and robots.[14]

Human-robot interaction refers to "how people work or play with robots."[15] Robots are quite different from other computers or tools because they are "physically situated agents," and human users interact with them in distinct ways.[16] Research on interaction needs to span a number of domains well beyond engineering, including psychology, cognitive science, and communications, among others.

"*Natural language processing* concerns … systems that can communicate with people using ordinary human languages."[17] Moreover, "natural language is the most normal and intuitive way for humans to instruct autonomous systems; it allows them to provide diverse, high-level goals and strategies rather than detailed teleoperation."[18] Hence, further development of the ability of autonomous weapons systems to respond to commands in a natural language is necessary.

Finally, the Defense Science Board uses the term *multiagent coordination* for circumstances in which a task is distributed among "multiple robots, software agents, or humans."[19] Tasks could be centrally planned or coordinated through interactions of the agents. This sort of coordination goes beyond mere cooperation because "it assumes that the agents have a cognitive understanding of each other's capabilities, can monitor progress towards the goal, and engage in more human-like teamwork."[20]

Moral justifications. Several military experts and roboticists have argued that autonomous weapons systems should not only be regarded as morally acceptable but also that they would in fact be ethically preferable

to human fighters. For example, roboticist Ronald C. Arkin believes autonomous robots in the future will be able to act more "humanely" on the battlefield for a number of reasons, including that they do not need to be programmed with a self-preservation instinct, potentially eliminating the need for a "shoot-first, ask questions later" attitude.[21] The judgments of autonomous weapons systems will not be clouded by emotions such as fear or hysteria, and the systems will be able to process much more incoming sensory information than humans without discarding or distorting it to fit preconceived notions. Finally, per Arkin, in teams comprised of human and robot soldiers, the robots could be more relied upon to report ethical infractions they observed than would a team of humans who might close ranks.[22]

Lt. Col. Douglas A. Pryer, U.S. Army, asserts there might be ethical advantages to removing humans from high-stress combat zones in favor of robots. He points to neuroscience research that suggests the neural circuits responsible for conscious self-control can shut down when overloaded with stress, leading to sexual assaults and other crimes that soldiers would otherwise be less likely to commit. However, Pryer sets aside the question of whether or not waging war via robots is ethical in the abstract. Instead, he suggests that because it sparks so much moral outrage among the populations from whom the United States most needs support, robot warfare has serious strategic disadvantages, and it fuels the cycle of perpetual warfare.[23]

Arguments Opposed to Autonomous Weapons Systems

While some support autonomous weapons systems with moral arguments, others base their opposition on moral grounds. Still others assert that moral arguments against autonomous weapons systems are misguided.

Opposition on moral grounds. In July 2015, an open letter calling for a ban on autonomous weapons was released at an international joint conference on artificial intelligence. The letter warns, "Artificial Intelligence (AI) technology has reached a point where the deployment of such systems is—practically if not legally—feasible within years, not decades, and the stakes are high: autonomous weapons have been described as the third revolution in warfare, after gunpowder and nuclear arms."[24] The letter also notes that AI has the potential to benefit humanity, but that if a military AI arms race ensues, AI's reputation

could be tarnished, and a public backlash might curtail future benefits of AI. The letter has an impressive list of signatories, including Elon Musk (inventor and founder of Tesla), Steve Wozniak (cofounder of Apple), physicist Stephen Hawking (University of Cambridge), and Noam Chomsky (Massachusetts Institute of Technology), among others. Over three thousand AI and robotics researchers have also signed the letter. The open letter simply calls for "a ban on offensive autonomous weapons beyond meaningful human control."[25]

We note in passing that it is often unclear whether a weapon is offensive or defensive. Thus, many assume that an effective missile defense shield is strictly defensive, but it can be extremely destabilizing if it allows one nation to launch a nuclear strike against another without fear of retaliation.

In April 2013, the United Nations (UN) special rapporteur on extrajudicial, summary, or arbitrary executions presented a report to the UN Human Rights Council. The report recommended that member states should declare and implement moratoria on the testing, production, transfer, and deployment of lethal autonomous robotics (LARs) until an internationally agreed upon framework for LARs has been established.[26]

That same year, a group of engineers, AI and robotics experts, and other scientists and researchers from thirty-seven countries issued the "Scientists' Call to Ban Autonomous Lethal Robots." The statement notes the lack of scientific evidence that robots could, in the future, have "the functionality required for accurate target identification, situational awareness, or decisions regarding the proportional use of force."[27] Hence, they may cause a high level of collateral damage. The statement ends by insisting that "decisions about the application of violent force must not be delegated to machines."[28]

Indeed, the delegation of life-or-death decision making to nonhuman agents is a recurring concern of those who oppose autonomous weapons systems. The most obvious manifestation of this concern relates to systems that are capable of choosing their own targets. Thus, highly regarded computer scientist Noel Sharkey has called for a ban on "lethal autonomous targeting" because it violates the Principle of Distinction, considered one of the most important rules of armed conflict—autonomous weapons systems will find it very hard to determine who is a civilian and who is a combatant, which is difficult

even for humans.[29] Allowing AI to make decisions about targeting will most likely result in civilian casualties and unacceptable collateral damage.

Another major concern is the problem of accountability when autonomous weapons systems are deployed. Ethicist Robert Sparrow highlights this ethical issue by noting that a fundamental condition of international humanitarian law, or *jus in bello*, requires that some person must be held responsible for civilian deaths. Any weapon or other means of war that makes it impossible to identify responsibility for the casualties it causes does not meet the requirements of jus in bello, and, therefore, should not be employed in war.[30]

This issue arises because AI-equipped machines make decisions on their own, so it is difficult to determine whether a flawed decision is due to flaws in the program or in the autonomous deliberations of the AI-equipped (so-called smart) machines. The nature of this problem was highlighted when a driverless car violated the speed limits by moving too slowly on a highway, and it was unclear to whom the ticket should be issued.[31] In situations where a human being makes the decision to use force against a target, there is a clear chain of accountability, stretching from whoever actually "pulled the trigger" to the commander who gave the order. In the case of autonomous weapons systems, no such clarity exists. It is unclear who or what are to be blamed or held liable.

Amitai Etzioni is a professor of international relations at The George Washington University. He served as a senior advisor at the Carter White House and taught at Columbia University, Harvard Business School, and the University of California at Berkeley. A study by Richard Posner ranked him among the top one hundred American intellectuals. His most recent book is *Foreign Policy: Thinking Outside the Box* (2016).

Oren Etzioni is chief executive officer of the Allen Institute for Artificial Intelligence. He received a PhD from Carnegie Mellon University and a BA from Harvard University. He has been a professor at the University of Washington's computer science department since 1991. He was the founder or cofounder of several companies, including Farecast (later sold to Microsoft) and Decide (later sold to eBay), and the author of over one hundred technical papers that have garnered over twenty-five thousand citations.

What Sharkey, Sparrow and the signatories of the open letter propose could be labeled "upstream regulation," that is, a proposal for setting limits on the development of autonomous weapons systems technology and drawing red lines that future technological developments should not be allowed to cross. This kind of upstream approach tries to foresee the direction of technological development and preempt the dangers such developments would pose. Others prefer "downstream regulation," which takes a wait-and-see approach by developing regulations as new advances occur. Legal scholars Kenneth Anderson and Matthew Waxman, who advocate this approach, argue that regulation will have to emerge along with the technology because they believe that morality will coevolve with technological development.[32]

Thus, arguments about the irreplaceability of human conscience and moral judgment may have to be revisited.[33] In addition, they suggest that as humans become more accustomed to machines performing functions with life-or-death implications or consequences (such as driving cars or performing surgeries), humans will most likely become more comfortable with AI technology's incorporation into weaponry. Thus, Anderson and Waxman propose what might be considered a communitarian solution by suggesting that the United States should work on developing norms and principles (rather than binding legal rules) guiding and constraining research and development—and eventual deployment—of autonomous weapons systems. Those norms could help establish expectations about legally or ethically appropriate conduct. Anderson and Waxman write,

> To be successful, the United States government would have to resist two extreme instincts. It would have to resist its own instincts to hunker down behind secrecy and avoid discussing and defending even guiding principles. It would also have to refuse to cede the moral high ground to critics of autonomous lethal systems, opponents demanding some grand international treaty or multilateral regime to regulate or even prohibit them.[34]

Counterarguments. In response, some argue against any attempt to apply to robots the language of morality that applies to human agents. Military ethicist George Lucas Jr. points out, for example, that robots cannot feel anger or a desire to "get even" by seeking retaliation for harm done to their compatriots.[35] Lucas holds that the debate thus far has been obfuscated by the confusion of machine autonomy with moral autonomy. The Roomba vacuum cleaner and Patriot missile "are both 'autonomous' in that they perform their assigned missions, including encountering and responding to obstacles, problems, and unforeseen circumstances with minimal human oversight," but not in the sense that they can change or abort their mission if they have "moral objections."[36] Lucas thus holds that the primary concern of engineers and designers developing autonomous weapons systems should not be *ethics* but rather *safety* and *reliability*, which means taking due care to address the possible risks of malfunctions, mistakes, or misuse that autonomous weapons systems will present. We note, though, that safety is of course a moral value as well.

Lt. Col. Shane R. Reeves and Maj. William J. Johnson, judge advocates in the U.S. Army, note that there are battlefields absent of civilians, such as underwater and space, where autonomous weapons could reduce the possibility of suffering and death by eliminating the need for combatants.[37] We note that this valid observation does not agitate against a ban in other, in effect most, battlefields.

Michael N. Schmitt of the Naval War College makes a distinction between weapons that are illegal per se and the unlawful use of otherwise legal weapons. For example, a rifle is not prohibited under international law, but using it to shoot civilians would constitute an unlawful use. On the other hand, some weapons (e.g., biological weapons) are unlawful per se, even when used only against combatants. Thus, Schmitt grants that some autonomous weapons systems might contravene international law, but "it is categorically not the case that all such systems will do so."[38] Thus, even an autonomous system that is incapable of distinguishing between civilians and combatants should not necessarily be unlawful per se, as autonomous weapons systems could be used in situations where no civilians were present, such as against tank formations in the desert or against warships. Such a system could be *used* unlawfully, though, if it were employed in contexts where civilians were present. We assert that some limitations on such weapons should be called for.

Soldiers from 2nd Battalion, 27th Infantry Regiment, 3rd Brigade Combat Team, 25th Infantry Division, move forward toward simulated opposing forces with a multipurpose unmanned tactical transport 22 July 2016 during the Pacific Manned-Unmanned Initiative at Marine Corps Training Area Bellows, Hawaii. (Photo by Staff Sgt. Christopher Hubenthal, U.S. Air Force)

In their review of the debate, legal scholars Gregory Noone and Diana Noone conclude that everyone is in agreement that any autonomous weapons system would have to comply with the Law of Armed Conflict (LOAC), and thus be able to distinguish between combatants and noncombatants. They write, "No academic or practitioner is stating anything to the contrary; therefore, this part of any argument from either side must be ignored as a red herring. Simply put, no one would agree to any weapon that ignores LOAC obligations."[39]

Limits on Autonomous Weapons Systems and Definitions of Autonomy

The international community has agreed to limits on mines and chemical and biological weapons, but an agreement on limiting autonomous weapons systems would meet numerous challenges. One challenge is the lack of consensus on how to define the autonomy of weapons systems, even among members of the Department of Defense. A standard definition that accounts for levels of autonomy could help guide an incremental approach to proposing limits.

Limits on autonomous weapons systems. We take it for granted that no nation would agree to forswear the use of autonomous weapons systems unless its adversaries would do the same. At first blush, it may seem that it is not beyond the realm of possibility to obtain an international agreement to ban autonomous weapons systems or at least some kinds of them.

Many bans exist in one category or another of weapons, and they have been quite well honored and enforced. These include the Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines and on their Destruction (known as the Ottawa Treaty, which became international law in 1999); the Chemical Weapons Convention (ratified in 1997); and the Convention on the Prohibition of the Development,

Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction (known as the Biological Weapons Convention, adopted in 1975). The record of the Treaty on the Nonproliferation of Nuclear Weapons (adopted in 1970) is more complicated, but it is credited with having stopped several nations from developing nuclear arms and causing at least one to give them up.

Some advocates of a ban on autonomous weapons systems seek to ban not merely production and de-

The *U.S. Army Robotic and Autonomous Systems Strategy*, published March 2017 by U.S. Army Training and Doctrine Command, describes how the Army intends to integrate new technologies into future organizations to help ensure overmatch against increasingly capable enemies. Five capability objectives are to increase situational awareness, lighten soldiers' workloads, sustain the force, facilitate movement and maneuver, and protect the force. To view the strategy, visit https://www.tradoc.army.mil/FrontPageContent/Docs/RAS_Strategy.pdf.

ployment but also research, development, and testing of these machines. This may well be impossible as autonomous weapons systems can be developed and tested in small workshops and do not leave a trail. Nor could one rely on satellites for inspection data for the same reasons. We hence assume that if such a ban were possible, it would mainly focus on deployment and mass production.

Even so, such a ban would face considerable difficulties. While it is possible to determine what is a chemical weapon and what is not (despite some disagreements at the margin, for example, about law enforcement use of irritant chemical weapons), and to clearly define nuclear arms or land mines, autonomous weapons systems come with very different levels of autonomy.[40] A ban on all autonomous weapons would require foregoing many modern weapons already mass produced and deployed.

Definitions of autonomy. Different definitions have been attached to the word "autonomy" in different Department of Defense documents, and the resulting concepts suggest rather different views on the future of robotic warfare. One definition, used by the Defense Science Board, views autonomy merely as high-end automation: "a capability (or a set of capabilities) that enables a particular action of a system to be automatic or, within programmed boundaries, 'self-governing.'"[41] According to this definition, already existing capabilities, such as autopilot used in aircraft, could qualify as autonomous.
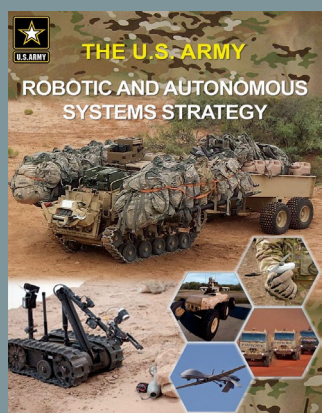
Another definition, used in the *Unmanned Systems Integrated Roadmap FY2011–2036*, suggests a qualitatively different view of autonomy: "an autonomous system is able to make a decision based on a set of rules and/or limitations. It is able to determine what information is important in making a decision."[42] In this view, autonomous systems are less predictable than merely automated ones, as the AI not only is performing a specified action but also is making decisions and thus potentially taking an action that a human did not order. A human is still responsible for programming the behavior of the autonomous system, and the actions the system takes would have to be consistent with the laws and strategies provided by humans. However, no individual action would be completely predictable or preprogrammed.

It is easy to find still other definitions of autonomy. The International Committee of the Red Cross defines autonomous weapons as those able to "independently select and attack targets, i.e., with autonomy in the 'critical functions' of acquiring, tracking, selecting and attacking targets."[43]

A 2012 Human Rights Watch report by Bonnie Docherty, *Losing Humanity: The Case against Killer Robots*, defines three categories of autonomy. Based on the kind of human involvement, the categories are human-in-the-loop, human-on-the-loop, and human-out-of-the-loop weapons.[44]

"Human-*in-the*-loop weapons [are] robots that can select targets and deliver force only with a human command."[45] Numerous examples of the first type already are in use. For example, Israel's Iron Dome system detects incoming rockets, predicts their trajectory, and then sends

this information to a human soldier who decides whether to launch an interceptor rocket.[46]

"Human-*on-the*-loop weapons [are] robots that can select targets and deliver force under the oversight of a human operator who can override the robots' actions."[47] An example mentioned by Docherty includes the SGR-A1 built by Samsung, a sentry robot used along the Korean Demilitarized Zone. It uses a low-light camera and pattern-recognition software to detect intruders and then issues a verbal warning. If the intruder does not surrender, the robot has a machine gun that can be fired remotely by a soldier the robot has alerted, or by the robot itself if it is in fully automatic mode.[48]

The United States also deploys human-on-the-loop weapons systems. For example, the MK 15 Phalanx Close-In Weapons System has been used on Navy ships since the 1980s, and it is capable of detecting, evaluating, tracking, engaging, and using force against antiship missiles and high-speed aircraft threats without any human commands.[49] The Center for a New American Security published a white paper that estimated as of 2015 at least thirty countries have deployed or are developing human-supervised systems.[50]

"Human-*out-of*-the-loop weapons [are] robots capable of selecting targets and delivering force without any human input or interaction."[51] This kind of autonomous weapons system is the source of much concern about "killing machines." Military strategist Thomas K. Adams warned that, in the future, humans would be reduced to making only initial policy decisions about war, and they would have mere symbolic authority over automated systems.[52] In the Human Rights Watch report, Docherty warns, "By eliminating human involvement in the decision to use lethal force in armed conflict, fully autonomous weapons would undermine other, nonlegal protections for civilians."[53] For example, a repressive dictator could deploy emotionless robots to kill and instill fear among a population without having to worry about soldiers who might empathize with their victims (who might be neighbors, acquaintances, or even family members) and then turn against the dictator.

For the purposes of this paper, we take autonomy to mean a machine has the ability to make decisions based on information gathered by the machine and to act on the basis of its own deliberations, beyond the instructions and parameters its producers, programmers, and users provided to the machine.

A Way to Initiate an International Agreement Limiting Autonomous Weapons

We find it hard to imagine nations agreeing to return to a world in which weapons had no measure of autonomy. On the contrary, development in AI leads one to expect that more and more machines and instruments of all kinds will become more autonomous. Bombers and fighter aircraft having no human pilot seem inevitable. Although it is true that any level of autonomy entails, by definition, some loss of human control, this genie has left the bottle and we see no way to put it back again.

Where to begin. The most promising way to proceed is to determine whether one can obtain an international agreement to ban *fully autonomous* weapons with missions that *cannot be aborted* and that cannot be recalled once they are launched. If they malfunction and target civilian centers, there is no way to stop them. Like unexploded landmines placed without marks, these weapons will continue to kill even after the sides settle their difference and sue for peace.

One may argue that gaining such an agreement should not be arduous because no rational policy maker will favor such a weapon. Indeed, the Pentagon has directed that "autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force."[54]

Why to begin. However, one should note that human-out-of-the-loop arms are very effective in reinforcing a red line. Declaration by representatives of one nation that if another nation engages in a certain kind of hostile behavior, swift and severe retaliation will follow, are open to misinterpretation by the other side, even if backed up with deployment of troops or other military assets.

Leaders, drawing on considerable historical experience, may bet that they be able to cross the red line and be spared because of one reason or another. Hence, arms without a human in the loop make for much more credible red lines. (This is a form of the "precommitment strategy" discussed by Thomas Schelling in *Arms and Influence*, in which one party limits its own

options by obligating itself to retaliate, thus making its deterrence more credible.)[55]

We suggest that nations might be willing to forgo this advantage of fully autonomous arms in order to gain the assurance that once hostilities ceased, they could avoid becoming entangled in new rounds of fighting because some bombers were still running loose and attacking the other side, or because some bombers might malfunction and attack civilian centers. Finally, if a ban on fully autonomous weapons were agreed upon and means of verification were developed, one could aspire to move toward limiting weapons with a high but not full measure of autonomy. ■

The authors are indebted to David Kroeker Maus for substantial research on this article.

Notes

1. Gary E. Marchant et al., "International Governance of Autonomous Military Robots," *Columbia Science and Technology Law Review* 12 (June 2011): 272–76, accessed 27 March 2017, http://stlr.org/download/volumes/volume12/marchant.pdf.

2. James R. Clapper Jr. et al., *Unmanned Systems Roadmap: 2007-2032* (Washington, DC: Department of Defense [DOD], 2007), 19, accessed 28 March 2017, http://www.globalsecurity.org/intell/library/reports/2007/dod-unmanned-systems-road-map_2007-2032.pdf.

3. Jeffrey S. Thurnher, "Legal Implications of Fully Autonomous Targeting," *Joint Force Quarterly* 67 (4th Quarter, October 2012): 83, accessed 8 March 2017, http://ndupress.ndu.edu/Portals/68/Documents/jfq/jfq-67/JFQ-67_77-84_Thurnher.pdf.

4. David Francis, "How a New Army of Robots Can Cut the Defense Budget," *Fiscal Times,* 2 April 2013, accessed 8 March 2017, http://www.thefiscaltimes.com/Articles/2013/04/02/How-a-New-Army-of-Robots-Can-Cut-the-Defense-Budget. Francis attributes the $850,000 cost estimate to an unnamed DOD source, presumed from 2012 or 2013.

5. Ibid.

6. Quoted in Evan Ackerman, "U.S. Army Considers Replacing Thousands of Soldiers with Robots," *IEEE Spectrum*, 22 January 2014, accessed 28 March 2016, http://spectrum.ieee.org/automaton/robotics/military-robots/army-considers-replacing-thousands-of-soldiers-with-robots.

7. Jason S. DeSon, "Automating the Right Stuff? The Hidden Ramifications of Ensuring Autonomous Aerial Weapon Systems Comply with International Humanitarian Law," *Air Force Law Review* 72 (2015): 85–122, accessed 27 March 2017, http://www.afjag.af.mil/Portals/77/documents/AFD-150721-006.pdf.

8. Michael Byrnes, "Nightfall: Machine Autonomy in Air-to-Air Combat," *Air & Space Power Journal* 23, no. 3 (May-June 2014): 54, accessed 8 March 2017, http://www.au.af.mil/au/afri/aspj/digital/pdf/articles/2014-May-Jun/F-Byrnes.pdf?source=GovD.

9. Defense Science Board, *Task Force Report: The Role of Autonomy in DoD Systems* (Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, July 2012), 31.

10. Ibid., 33.

11. Ibid., 38–39.

12. Ibid., 39.

13. Ibid., 41.

14. Ibid., 42.

15. Ibid., 44.

16. Ibid.

17. Ibid., 49.

18. Ibid.

19. Ibid., 50.

20. Ibid.

21. Ronald C. Arkin, "The Case for Ethical Autonomy in Unmanned Systems," *Journal of Military Ethics* 9, no. 4 (2010): 332–41.

22. Ibid.

23. Douglas A. Pryer, "The Rise of the Machines: Why Increasingly 'Perfect' Weapons Help Perpetuate Our Wars and Endanger Our Nation," *Military Review* 93, no. 2 (2013): 14–24.

24. "Autonomous Weapons: An Open Letter from AI [Artificial Intelligence] & Robotics Researchers," Future of Life Institute website, 28 July 2015, accessed 8 March 2017, http://futureoflife.org/open-letter-autonomous-weapons/.

25. Ibid.

26. *Report of the Special Rapporteur on Extrajudicial, Summary, or Arbitrary Executions, Christof Heyns*, September 2013, United Nations Human Rights Council, 23rd Session, Agenda Item 3, United Nations Document A/HRC/23/47.

27. International Committee for Robot Arms Control (ICRAC), "Scientists' Call to Ban Autonomous Lethal Robots," ICRAC website, October 2013, accessed 24 March 2017, icrac.net.

28. Ibid.

29. Noel Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9, no. 4 (2010): 369–83, accessed 28 March 2017, doi:10.1080/15027570.2010.537903. For more on this subject, see Peter Asaro, "On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-making," *International Review of the Red Cross* 94, no. 886 (2012): 687–709.

30. Robert Sparrow, "Killer Robots," *Journal of Applied Philosophy* 24, no. 1 (2007): 62–77.

31. For more discussion on this topic, see Amitai Etzioni and Oren Etzioni, "Keeping AI Legal," *Vanderbilt Journal of Entertainment & Technology Law* 19, no. 1 (Fall 2016): 133–46, accessed 8 March 2017, http://www.jetlaw.org/wp-content/uploads/2016/12/Etzioni_Final.pdf.

32. Kenneth Anderson and Matthew C. Waxman, "Law and Ethics for Autonomous Weapon Systems: Why a Ban Won't Work and How the Laws of War Can," Stanford University, Hoover Institution Press, Jean Perkins Task Force on National Security and Law Essay Series, 9 April 2013.

33. Ibid.

34. Anderson and Waxman, "Law and Ethics for Robot Soldiers," *Policy Review* 176 (December 2012): 46.

35. George Lucas Jr., "Engineering, Ethics & Industry: the Moral Challenges of Lethal Autonomy," in *Killing by Remote Control: The Ethics of an Unmanned Military*, ed. Bradley Jay Strawser (New York: Oxford, 2013).

36. Ibid., 218.

37. Shane Reeves and William Johnson, "Autonomous Weapons: Are You Sure these Are Killer Robots? Can We Talk About It?," in Department of the Army Pamphlet 27-50-491, *The Army Lawyer* (Charlottesville, VA: Judge Advocate General's Legal Center and School, April 2014), 25–31.

38. Michael N. Schmitt, "Autonomous Weapon Systems and International Humanitarian Law: a Reply to the Critics," *Harvard National Security Journal*, 5 February 2013, accessed 28 March 2017, http://harvardnsj.org/2013/02/autonomous-weapon-systems-and-international-humanitarian-law-a-reply-to-the-critics/.

39. Gregory P. Noone and Diana C. Noone, "The Debate over Autonomous Weapons Systems," *Case Western Reserve Journal of International Law* 47, no. 1 (Spring 2015): 29, accessed 27 March 2017, http://scholarlycommons.law.case.edu/jil/vol47/iss1/6/.

40. Neil Davison, ed., *'Non-lethal' Weapons* (Houndmills, England: Palgrave Macmillan, 2009).

41. DOD Defense Science Board, *Task Force Report: The Role of Autonomy in DOD Systems*, 1.

42. DOD, *Unmanned Systems Integrated Roadmap FY2011-2036* (Washington, DC: Government Publishing Office [GPO], 2011), 43.

43. International Committee of the Red Cross (ICRC), Expert Meeting 26–28 March 2014 report, "Autonomous Weapon Systems: Technical, Military, Legal and Humanitarian Aspects" (Geneva: ICRC, November 2014), 5.

44. Bonnie Docherty, *Losing Humanity: The Case against Killer Robots* (Cambridge, MA: Human Rights Watch, 19 November 2012), 2, accessed 10 March 2017, https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots.

45. Ibid.

46. Paul Marks, "Iron Dome Rocket Smasher Set to Change Gaza Conflict," New Scientist Daily News online, 20 November 2012, accessed 24 March 2017, https://www.newscientist.com/article/dn22518-iron-dome-rocket-smasher-set-to-change-gaza-conflict/.

47. Docherty, *Losing Humanity,* 2.

48. Ibid.; Patrick Lin, George Bekey, and Keith Abney, *Autonomous Military Robotics: Risk, Ethics, and Design* (Arlington, VA: Department of the Navy, Office of Naval Research, 20 December 2008), accessed 24 March 2017, http://digitalcommons.calpoly.edu/cgi/viewcontent.cgi?article=1001&context=phil_fac.

49. "MK 15—Phalanx Close-In Weapons System (CIWS)" Navy Fact Sheet, 25 January 2017, accessed 10 March 2017, http://www.navy.mil/navydata/fact_print.asp?cid=2100&tid=487&ct=2&page=1.

50. Paul Scharre and Michael Horowitz, "An Introduction to Autonomy in Weapons Systems" (working paper, Center for a New American Security, February 2015), 18, accessed 24 March 2017, http://www.cnas.org/.

51. Docherty, *Losing Humanity*, 2.

52. Thomas K. Adams, "Future Warfare and the Decline of Human Decisionmaking," *Parameters* 31, no. 4 (Winter 2001–2002): 57–71.

53. Docherty, *Losing Humanity*, 4.

54. DOD Directive 3000.09, *Autonomy in Weapon Systems* (Washington, DC: U.S. GPO, 21 November 2012), 2, accessed 10 March 2017, http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf.

55. Thomas C. Schelling, *Arms and Influence* (New Haven: Yale University, 1966).

Office of the Secretary of Defense Logistics Program fellows visit the bridge of the USNS *Gilliland* 28 October 2015 during a tour of the Military Sealift Command (MSC) ship. Shown from left to right are Stanley McMillian, Lt. Col. Ed Hogan (*kneeling*), Bryan Jerkatis, Donald Gillespie, Art Clark (MSC Surge Sealift Program Readiness Officer), USNS *Gilliland*'s Master Keith Finnerty, and Renee Hubbard. (Photo courtesy of U.S. Navy)

The Office of the Secretary of Defense Logistics Fellows Program

The Office of the Secretary of Defense (OSD) Logistics Fellows Program is open to field-grade officer (O4 to O5) and Department of Defense (DOD) civilian logisticians (GS-13 to GS-14). According to the Office of the Assistant Secretary of Defense for Logistics and Materiel Readiness, which administers the program, the one-year fellowship is a developmental assignment that fosters learning, growth, and experiential opportunities.

Fellows participate in policy formulation and department-wide oversight responsibilities. They tour public- and private-sector logistics organizations to learn how they conduct operations and to observe industry best practices. They gain insight into legislative processes through visits to Congress, and they attend national-level forums and engage in collaborative efforts with industry partners. Fellows can observe and interact with appointed and career senior executives and flag officers, including one-on-one meetings with senior logistics leaders in the military departments, joint staff, OSD, and other agencies. They can gain a deeper understanding of the OSD perspective and how it affects the DOD enterprise. More information about the program and how to apply can be found at http://www.acq.osd.mil/log/lmr/fellows_program.html.