

Chapter 16

Pros and Cons of Autonomous Weapons Systems (with Oren Etzioni)

Autonomous weapons systems (AWS) and military robots are progressing from science fiction movies to designer's 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 are able to perform increasingly advanced functions, including targeting and application of force, with little or no human oversight. Some military experts hold that these autonomous weapons systems not only confer significant strategic and tactical advantages in the battleground, but that they are also preferable to the use of human combatants, on moral grounds. In contrast, critics hold that these weapons should be curbed, if not banned altogether, for a variety of moral and legal reasons. The chapter reviews first the arguments by those who favor AWS, then those who oppose them, and closes with a policy suggestion.

16.1 In Support of AWS

16.1.1 *Military Advantages*

Those who call for further development and deployment of autonomous weapons systems generally point to several advantages. (a) Autonomous weapons systems act as a “force multiplier;” that is, fewer soldiers are needed for a given mission, and the efficacy of each soldier is greater. (b) Autonomous weapons systems expand the battlefield, allowing combat to reach into areas that were previously inaccessible. And (c) Autonomous weapons systems reduce casualties by removing human soldiers from dangerous missions (Marchant et al. 2011, pp. 272–276).

This chapter draws on “Pros and Cons of Autonomous Weapons Systems” in *Military Review*, (May-June 2017): 72–80.

The Pentagon's *Unmanned Systems Roadmap 2007–2032* provides additional motivations for pursuing AWS. These include that robots are better suited than humans for “dull,” “dangerous,” and “dirty” missions. Examples given for each respective category of mission include long sorties, bomb disposal, and operating in nuclear clouds or areas with high radioactivity (Clapper et al. 2007). Jeffrey S. Thurnher of the US Naval War College adds that “LARs [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” (Thurnher 2012, p. 83).

The long-term savings that could be achieved through fielding an army of military robots have also been highlighted. *The Fiscal Times* notes that each US soldier in Afghanistan costs the Pentagon roughly \$850,000 per year (some estimate the cost to be over \$1 million per soldier per year), which does not include the long-term costs of providing health care to veterans. Conversely, the TALON robot—a small, armed robot—can be built for only \$230,000 and is relatively cheap to maintain (Francis 2013). Gen. Robert Cone, head of the Army's Training and Doctrine Command, suggested in 2014 that by relying more on “support robots,” the Army could reduce the size of a brigade from 4000 to 3000 soldiers without a concomitant reduction in effectiveness (Ackerman 2014).

Major Jason DeSon, writing in the *Air Force Law Review*, notes the potential advantages of autonomous aerial weapons systems. The physical strain of high-g maneuvers and the intense mental concentration and situational awareness required of fighter pilots makes 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, which could confuse an opponent (DeSon 2015). More striking still, US Air Force Captain Michael Byrnes predicts that a single Unmanned Aerial Vehicle (UAV) 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 with human pilots (Byrnes 2014).

In guiding future research in AWS, the Defense Science Board at the Pentagon has identified six areas where advances in autonomy would be of significant benefit to current systems:

Perception, which includes not just new hardware (the actual sensors) but also software (algorithms for sensing).

Planning, which includes “the algorithms needed to make decisions about action (provide autonomy) in situations in which humans are not in the environment (e.g. space, the ocean)” (DSB 2012, p. 39).

Learning. The DSB report states the advantages of machine learning over manual software engineering, but notes that machine learning approaches to autonomous vehicles have thus far mostly been applied to ground vehicles and robots, and not yet air and marine vehicles.

Human-Robot Interaction (HRI). Robots are quite different from other computers or tools because they are physically situated agents, and thus elicit different

responses from human users. Hence HRI research needs to span a number of domains well beyond engineering, including psychology, cognitive science, and communications, among others.

Natural language. The authors of the DSB report hold that “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” (DSB 2012, p. 49). Hence, further development of the ability of autonomous weapons systems to respond to commands in a natural language is necessary.

Multi-Agent Coordination refers to the distribution of tasks among multiple robots, with either centrally planned or directly negotiated synchronization. This sort of collaboration 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” (DSB 2012, p. 50).

16.1.2 Moral Justifications

Several military experts and roboticists have argued that autonomous weapons systems should not only be regarded as morally acceptable, but that they would in fact be ethically preferable to human fighters. Roboticist Ronald Arkin believes that autonomous robots in the future will be able to act more “humanely” on the battlefield for a number of reasons: For one, they do not need to be programmed with a self-preservation instinct, thus potentially eliminating the need for a “shoot-first, ask questions later” attitude. The judgments of autonomous weapons systems will not be clouded by emotions like fear or hysteria, and they will be able to process much more incoming sensory information than humans, without discarding or distorting it to fit preconceived notions. Finally, in teams comprised of human and robot soldiers, the robots could be more relied upon to report ethical infractions that they observe than would a team of humans who might close ranks (Arkin 2010).

Lieutenant Colonel Douglas A. Pryer of the US Army adds that there might be ethical advantages to removing humans from high-stress combat zones in favor of robots. He points to neuroscience research which suggests that 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. But Pryer (2013) sets aside the question of whether or not waging war via robots is ethical in the abstract and suggests that because it sparks so much moral outrage among the populations from which the US most needs support, robot warfare has serious strategic disadvantages and is helping to fuel the cycle of perpetual warfare.

16.2 Opposition to AWS

16.2.1 *Opposition on Moral Grounds*

In July of 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” (Autonomous Weapons 2015). The letter also notes that AI has the potential to benefit humanity, but that if a military AI arms race ensues, its 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 (co-founder of Apple), physicist Stephen Hawking (University of Cambridge), and Noam Chomsky (MIT), among others. Over 3000 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.” 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 defensive, but it can be extremely destabilizing, if it allows one nation to launch a nuclear strike against another without fear of retaliation.

Previously, in April of 2013, the UN’s Special Rapporteur on extrajudicial, summary, and arbitrary executions presented a report to the UN’s Human Rights Council recommending 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 (Heyns 2013).

That same year, a group of engineers, AI and robotics experts, and other scientists and researchers from 37 countries issued the “Scientists’ Call to Ban Autonomous Lethal Robots.” The statement notes the lack of scientific evidence that robots could, in the foreseeable future, have “the functionality required for accurate target identification, situational awareness or decisions regarding the proportional use of force.” 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” (ICRAC 2013).

Indeed, the delegation of life-or-death decision-making to non-human agents is a recurring concern of AWS’ opponents. The most obvious manifestation of this concern relates to autonomous weapons systems that are capable of choosing their own targets. Thus, highly regarded computer scientist Noel Sharkey (2012) has called for a ban on “autonomous lethal 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. Allowing AI to make decisions about targeting will most likely result in civilian casualties and unacceptable collateral damage.

Another major concern deals with the problem of accountability when autonomous weapons systems are deployed. Ethicist Robert Sparrow (2007) highlights this ethical issue by noting that a fundamental condition of international humanitarian law, or *jus in bello*, requires that someone must be able to 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.

This issue arises because AI-equipped machines make decisions on their own, which makes it 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. 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 (For more, see Etzioni and Eztioni 2016). 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 orders. In the case of AWS, no such clarity exists. It is unclear who or what is to blame or bears liability.

What Sharkey, Sparrow, and the signatories of the open letter propose could be labelled “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 pre-empt 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 co-evolve with technological development. Thus, arguments about the irreplaceability of human conscience and moral judgment may have to be revisited (Anderson and Waxman 2013a). They suggest that, as humans become more accustomed to machines performing functions with life-or-death implications/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 AWS. Those norms could help establish expectations about legally or ethically appropriate conduct. They 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. (Anderson and Waxman 2013b, p. 46)

16.2.2 Counter Arguments

In response, some argue against any attempt to apply the language of morality that is applied to human agents, to robots. Military ethicist George Lucas Jr. (2013) points out, for example, that robots cannot feel anger or a desire to “get even” by seeking retaliation for harm done to their compatriots. 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” (Lucas 2013, p. 218). 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.

Lieutenant Colonel Shane R. Reeves & Major William J. Johnson, Judge Advocates in the US Army, note that there are battlefields absent of civilians, such as underwater and in space, where autonomous weapons could reduce the possibility of suffering and death by eliminating the need for combatants (Reeves and Johnson 2014). 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” (Schmitt 2013, p. 8). 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 are present, such as against tank formations in the desert, or warships. Such a system could be *used* unlawfully though, if it were employed in contexts where civilians were present. We note that setting some limitations on such weapons should still be called for.

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” (Noone and Noone 2015, p. 25).

16.2.3 *Level of Autonomy*

We take it for granted that no nation will 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. One notes that a fair number of bans on one category or another of weapons exist and 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; the Chemical Weapons Convention; and the Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction. The record of the Treaty on the Non-Proliferation of Nuclear Weapons 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 of the advocates of a ban on autonomous weapons systems seek to ban not merely production and deployment but also R&D and testing of these machines. This may well not be possible 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 perhaps encompass 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 Davidson 2009), and to clearly define nuclear arms or land mines, autonomous weapons systems come with very different levels of autonomy. A ban on all autonomous weapons would require foregoing many modern weapons, already mass produced and deployed.

16.2.4 *Defining 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 Task Force 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’” (DSB 2012, p. 1). According to this definition, already existing capabilities, such as auto-pilot used in aircrafts, could qualify as “autonomous.”

Another definition, used in the DoD’s *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” (DOD 2011, p. 43). In this view, autonomous systems are less predictable than merely automated ones, as the AI is not only performing a specified action, but also 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, but 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” (ICRC 2014, p. 7).

We take autonomy to mean an ability to make decisions based on information gathered by a machine and to act on the basis of its own deliberations, beyond the instructions and parameters provided to the machine by its producers, programmers, and users.

It seems useful to consider three kinds or levels of autonomy:

Human-in-the-Loop Weapons: Robots that can select targets and deliver force only with a human command. Numerous examples of this type already exist and are in use. For example, Israel’s Iron Dome system detects incoming rockets, predicts its trajectory and then sends this information to a human soldier who decides whether to launch an interceptor rocket (Marks 2012).

Human-on-the-Loop Weapons: Robots that can select targets and deliver force under the oversight of a human operator who can override the robots’ actions. Examples of Human-on-the-Loop weapons are either in development or have already been deployed. The SGR-A1 built by Samsung is a sentry robot placed by South Korea along the Demilitarized Zone (DMZ). 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 which can either be fired remotely by a soldier, who the robot has alerted, or by the robot itself if it is in fully-automatic mode (Lin et al. 2008).

The US also deploys Human-on-the-Loop weapons systems. For example, the MK 15-Phalanx Close-In Weapons System has been used on US Navy ships since the 1980s and is capable of detecting, evaluating, tracking, engaging, and using force against Anti-Ship Missiles and high-speed aircraft threats without any human commands (MK 15 n.d.). The Center for a New American Security estimates that at least 30 countries have deployed or are developing Human-on-the-Loop systems (Scharre and Horowitz 2015).

Human-out-of-the-Loop Weapons: Robots that are capable of selecting targets and delivering force without any human input or interaction (Docherty 2012). This kind of autonomous weapons system is the source of much concern about “killing machines.” Military strategist Thomas K. Adams (2002) warned that, in the future, humans would be reduced to making only initial policy decisions about war and have mere symbolic authority over automated systems. Human Rights Watch, in its much discussed report *Losing Humanity: The Case against Killer*

Robots, warned, “By eliminating human involvement in the decision to use lethal force in armed conflict, fully autonomous weapons would undermine other, non-legal protections for civilians” (Docherty 2012). The authors believe that a repressive dictator could deploy emotionless robots to kill and instill fear among the population without having to worry about soldiers empathizing with their victims (who might be neighbors, acquaintances, or even family members) and turning against the dictator.

As we see it, it is hard to imagine nations agreeing to return to a world in which weapons have 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 with no human pilot seem inevitable. Although it is true that any degree 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.

The most promising way to proceed is to determine whether one can obtain international agreement to ban fully autonomous weapons, whose missions 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 make peace.

One may argue that gaining such an agreement should not be arduous because no rational policymaker 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” (DOD 2012). One should note though that such 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. Arms without a human in the loop make for much more credible red lines. (This is a form of the “pre-commitment strategy” discussed by Thomas Schelling (1966) in *Arms and Influence*, in which one party limits its own options by obligating itself to retaliate, thus making its deterrence more credible.)

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

References

- Ackerman, E. 2014. U.S. Army considers replacing thousands of soldiers with robots. *IEEE Spectrum*.
- Adams, T.K. 2002. Future warfare and the decline of human decisionmaking. *Parameters* 31 (4): 57–71.
- Anderson, K., and M. Waxman. 2013a. *Law and ethics for autonomous weapon systems: why a ban won't work and how the laws of war can*. Stanford University, Hoover Institution (Jean Perkins Task Force on National Security and Law Essay Series).
- . 2013b. Law and ethics for robot soldiers. *Policy Review*.
- Arkin, R.C. 2010. The case for ethical autonomy in unmanned systems. *Journal of Military Ethics* 9 (4): 332–341.
- Autonomous Weapons. 2015. An open letter from AI & robotics researchers. <http://futureoflife.org/open-letter-autonomous-weapons/>.
- Byrnes, M. 2014. Nightfall: Machine autonomy in air-to-air combat. *Air and Space Power Journal* 28 (3): 48–75.
- Clapper, J., J. Young, J. Cartwright, and J. Grimes. 2007. *Unmanned systems roadmap 2007–2032*. Office of the Secretary of Defense.
- Davison, N. 2009. *'Non-lethal' weapons*. New York: Palgrave Macmillan.
- DeSon, J.S. 2015. Automating the right stuff – the hidden ramifications of ensuring autonomous aerial weapon systems comply with international humanitarian law. *Air Force Law Review* 72: 85–122.
- Docherty, B. 2012. Losing humanity: The case against killer robots. *Human Rights Watch*.
- DOD (Department of Defense). 2011. Unmanned systems integrated roadmap FY2011–2036.
- . 2012. Directive number 3000.09, Subject: Autonomy in weapons systems.
- DSB (Defense Science Board). 2012. Task force report: The role of autonomy in DoD systems.
- Etzioni, A., and O. Etzioni. 2016. Keeping AI legal. *Vanderbilt Journal of Entertainment & Technology Law* 19 (1): 133–164.
- Francis, D. 2013. How a new army of robots can cut the defense budget. *The Fiscal Times*.
- Heyns, C. 2013. *Report of the special rapporteur on extrajudicial, summary or arbitrary executions*. U.N. Human Rights Council. U.N. Doc. A/HRC/23/47.
- ICRAC (International Committee for Robot Arms Control). 2013. *Scientists' call to ban autonomous lethal robots*. <https://icrac.net/call/>.
- ICRC (International Committee of the Red Cross). 2014. Autonomous weapons systems: technical, military, legal, and humanitarian aspects.
- Lin, P., G. Bekey, K. Abney. 2008. *Autonomous military robotics: Risk, ethics, and design* (prepared for U.S. Department of Navy, Office of Naval Research), http://digitalcommons.calpoly.edu/cgi/viewcontent.cgi?article=1001&context=phil_fac.
- Lucas, G.R., Jr. 2013. Engineering, ethics & industry: The moral challenges of lethal autonomy. In *Killing by remote control: The ethics of an unmanned military*, ed. B. Strawser, 211–228. Oxford: Oxford University Press.
- Marchant, G.E., B. Allenby, R. Arkin, E.T. Barrett, J. Borenstein, L.M. Gaudet, et al. 2011. International governance of autonomous military robots. *Columbia Science & Technology Law Review* 12: 272–315.
- Marks, P. 2012. Iron dome rocket smasher set to change Gaza conflict. *New Scientist*.
- MK 15 – Phalanx close-in weapons system (CIWS). n.d.. *U.S. Navy fact sheet*. http://www.navy.mil/navydata/fact_print.asp?cid=2100&tid=487&ct=2&page=1.
- Noone, G.P., and D.C. Noone. 2015. The debate over autonomous weapons systems. *Case Western Reserve Journal of International Law* 47 (1): 25–35.
- Pryer, D.A. 2013. The rise of the machines: Why increasingly “perfect” weapons help perpetuate our wars and endanger our nation. *Military Review* 93 (2): 14–24.
- Reeves, S., and W. Johnson. 2014. Autonomous weapons: Are you sure these are killer robots? Can we talk about it? *The Army Lawyer (April)*: 25–31.

- Scharre, P., and M. Horowitz. 2015. *An introduction to autonomy in weapons systems* (Working Paper). Center for New American Security.
- Schelling, T.C. 1966. *Arms and influence*. New Haven: Yale University Press.
- Schmitt, M. 2013. Autonomous weapon systems and international humanitarian law: A reply to the critics. *Harvard National Security Journal Features*.
- Sharkey, N. 2012. Saying 'no!' to lethal autonomous targeting. *Journal of Military Ethics* 9 (4): 369–383.
- Sparrow, R. 2007. Killer robots. *Journal of Applied Philosophy* 24 (1): 62–77.
- Thurnher, J. 2012. Legal implications of fully autonomous targeting. *Joint Force Quarterly* 67 (4): 77–84.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

