

Chapter 13

Examining Autonomous Weapon Systems from a Law of Armed Conflict Perspective

Jeffrey S. Thurnher

Abstract This chapter explores the legal implications of autonomous weapon systems and the potential challenges such systems might present to the laws governing weaponry and the conduct of hostilities. Autonomous weapon systems are weapons that are capable of selecting and engaging a target without further human operator involvement. Although such systems have not yet been fully developed, technological advances, particularly in artificial intelligence, make the appearance of such systems a distinct possibility in the years to come. Given such a possibility, it is essential to look closely at both the relevant technology involved in these cutting-edge systems and the applicable law. This chapter commences with an examination of the emerging technology supporting these sophisticated systems, by detailing autonomous features that are currently being designed for weapons and anticipating how technological advances might be incorporated into future weapon systems. A second aim of the chapter is to describe the relevant law of armed conflict principles applicable to new weapon systems, with a particular focus on the unique legal challenges posed by autonomous weapons. The legal analysis will outline how autonomous weapon systems would need to be designed for them to be deemed lawful *per se*, and whether the use of autonomous weapons during hostilities might be prohibited in particular circumstances under the law of armed conflict. The third and final focus of this chapter is to address potential lacunae in the law dealing with autonomous weapon systems. In particular, the author will reveal how interpretations of and issues related to subjectivity in targeting decisions and overall accountability may need to be viewed differently in response to autonomy.

The author is Lieutenant Colonel in the United States Army, Judge Advocate, Faculty, International Law Department, Naval War College, Newport, Rhode Island, USA. The views expressed are those of the author and should not be understood as necessarily representing those of the United States Department of Defense or any other government entity.

J. S. Thurnher (✉)

Faculty, International Law Department, Naval War College, Newport, RA, USA
e-mail: jeffrey.s.thurnher.mil@mail.mil

Contents

13.1	Introduction.....	214
13.2	The Context: Autonomous Weapon Systems Technology	215
	13.2.1 Current State and Uses of Autonomy.....	216
	13.2.2 Future Technological Possibilities	218
13.3	Legal Implications of Autonomous Weapon Systems.....	219
	13.3.1 Weapons Law Rules.....	219
	13.3.2 Targeting Law Requirements	220
13.4	Legal Lacunae in Autonomous Weapon Systems?.....	222
	13.4.1 Subjectivity in Targeting.....	223
	13.4.2 Responsibility.....	225
13.5	Conclusion	226
	References.....	227

13.1 Introduction

The aim of this chapter is to explore the myriad of law of armed conflict issues surrounding the potential future development and deployment of autonomous weapon systems (AWS). The United States (US) military defines AWS as those weapons that ‘once activated, can select and engage targets without further human intervention by a human operator’.¹ While many existing weapon systems possess features which are autonomous, fully autonomous weapons do not currently exist in any nation’s military arsenal. Experts predict this will change in the future as AWS begin to supplant manned and remotely piloted systems. Some projections are that within 20 years AWS will be the primary weapons used on the battlefield.² Fearful that these systems might not be able to comport with international law, several non-governmental organisations and human rights advocacy groups have already voiced their opposition to AWS and have called for a pre-emptive international treaty banning their development and use.³ Furthermore, in April 2013, a United Nations (UN) Special Rapporteur issued a report to the UN Human Rights Council recommending a moratorium on all AWS testing and deployments until nations can agree on a legal and regulatory framework for their use.⁴

States are keenly aware of the concerns over the systems. States recognise that uncertainties surround the potential deployment of AWS, and they are sensitive to

¹ US Department of Defense 2012a, p. 13.

² Singer 2009, p. 128.

³ See, for example, Human Rights Watch 2012, p. 1.

⁴ Heyns 2013.

the issues. Some states have begun implementing measures meant to ensure compliance with the law of armed conflict.⁵ Yet the draw of autonomous weapons is significant. AWS offer tremendous promise in terms of protecting one's own forces from harm, and they may even be able to provide greater protections to civilians by delivering more precise and accurate strikes. Given such possibilities, it seems unlikely at this stage that a consensus of states will emerge to preemptively ban these systems. This chapter critically examines the crux of this emerging debate, namely whether AWS can be designed and used in ways that comply with the fundamental principles of the law of armed conflict.

In order to fully detail the legal implications of AWS, the chapter is organised as follows. [Section 13.2](#) explains how militaries are presently using autonomous technology and predicts how autonomy may be embedded into weapons in the future. [Section 13.3](#) describes the relevant tenets of the law of armed conflict in terms of identifying whether an autonomous weapon itself would be lawful and whether such a weapon could be used in a lawful manner. [Section 13.4](#) then identifies areas of the law that may need to be evaluated in a new light in response to the unique challenges raised by AWS. The author will focus particularly on the subjective judgements inherent in targeting and on responsibility. Finally, [Sect. 13.5](#) draws conclusions as to whether AWS will ultimately be deemed lawful under the law of armed conflict.

13.2 The Context: Autonomous Weapon Systems Technology

Technology and science have dramatically advanced warfare and improved the capabilities of weapons throughout history, but the emergence of autonomous technology may well represent a revolution for modern warfare. Humans have traditionally always been 'in the loop' with regard to lethal targeting decisions. The potential creation of autonomous systems capable of independently selecting and engaging targets with lethal force may alter that paradigm. This shift may occur rapidly if recent leaps in artificial intelligence are any indication. Some experts even contend that the technology to develop fully autonomous weapons essentially exists today.⁶

⁵ The US promulgated a policy directive in late 2012 establishing a strict approval process for any AWS acquisitions or development and mandating that various safety measures be incorporated into future AWS designs: US Department of Defense [2012a](#).

⁶ For example, the former chief scientist for the US Air Force even contends that technology currently exists to facilitate 'fully autonomous military strikes': Dahm [2012](#), p. 11.

13.2.1 Current State and Uses of Autonomy

Tremendous strides taken by artificial intelligence researchers in recent years make the prospects for fully autonomous systems more certain. Computer scientists are now successfully coupling sophisticated computer algorithms together in novel approaches designed to allow wider separation between human operators and robotic systems.⁷ Innovative computing approaches, including those using so called ‘machine learning’ processes, have enabled computers to more closely simulate human thought patterns.⁸ Increasingly, these computer systems are able to decipher answers to complex problems by ‘improv[ing] automatically through experience,’ an approach which is similar to the way humans learn by example.⁹ Given these artificial intelligence enhancements, embedded autonomous capabilities are becoming more prevalent across society. For instance, automobile manufacturers are outfitting new vehicles with a host of autonomous features, and many car makers and experts now predict that self-driving vehicles will become the norm in coming decades.¹⁰ The uses of autonomous technology are not, however, confined to commercial products. Militaries around the world are quickly embracing these opportunities.

Militaries are particularly well situated to capitalise on these advances because they have consistently been at the forefront of innovation. Advanced militaries began including automatic or autonomous features into their weaponry years ago. The US Navy, for example, has long been using naval mines that respond automatically to ‘acoustic, electromagnetic or pressure’ signatures.¹¹ The Navy has also been using close-in weapon systems, like the Phalanx and other similar systems, on warships since the mid-1970s as a protective measure of last resort in self-defence. The systems are designed to defeat incoming missile and rocket attacks against ships by reacting automatically with lethal force in response to the signatures of such threats. Many air defence weapon systems, like the US Army’s Patriot Missile system and the Israeli Iron Dome, have also been able to function for many years with various degrees of autonomy to defeat incoming artillery or missile attacks on ground forces.

While the above examples showcase limited uses of autonomous technologies the newest defence systems are set to truly harness the latest innovations in artificial intelligence and represent a significant step towards the development of fully autonomous systems. One example is the ‘K-MAX’ variant helicopters developed by the US Army and Marines, which have already flown autonomously along pre-

⁷ Poitras 2012.

⁸ For an overview of machine learning capabilities and possibilities, see, Russell and Norvig 2010, Chap. 18.

⁹ Public Broadcasting Service 2011.

¹⁰ IEEE 2012.

¹¹ Heintschel von Heinegg 2011, p. 184 (stating that these mines are ‘quite common and legally uncontested’). See also, Chap. 14 by McLaughlin in this volume.

programmed routes in Afghanistan to deliver cargo to forward operating bases. The US Navy has also produced a combat aircraft, known as the X-47B, designed to autonomously take off and land on an aircraft carrier. The British Royal Air Force is developing the Taranis attack aircraft, which will be capable of supersonic autonomous flight. The US Defense Advanced Research Projects Agency (DARPA) has begun perfecting autonomous mid-air refueling techniques. The US has even developed underwater systems which are capable of autonomously adjusting themselves to maintain their position in the water for months. Although these exceptional developments in autonomy have not yet included autonomous attack features, that is expected to change as these systems become further refined.

Militaries will likely pursue autonomous targeting capabilities, in part, to counter several perceived operational gaps and shortcomings with the current fleet of manned and remotely controlled systems. First, current unmanned systems that are operated remotely by a human pilot are susceptible to the ever-increasing communications jamming and cyber attack capabilities of adversaries. If the communications link between the human controller and the unmanned system is cut, the system becomes unable to complete its assigned mission. An AWS, on the other hand, will likely fare better in such electronically contested environments as it is not dependent on a tethered link to an operator. Second, remotely controlled systems are heavily dependent on large numbers of pilots and analysts. As the demand for unmanned systems continues to grow (and correspondingly, the volume of data generated by these systems exponentially expands),¹² these support requirements may increase to the point of becoming prohibitively burdensome for militaries. Autonomous systems will require fewer human observers as tasks are instead delegated to computerised systems.¹³ Third, the pace of combat in the future is expected to become too fast for human operators. Manned and remotely piloted systems might simply prove too slow and ineffective against an enemy who possesses autonomous systems.¹⁴ Rather than being put at a competitive disadvantage, more nations will likely seek this capability and develop AWS.¹⁵ Given these operational realities, it is apparent that AWS will play a continuing and more expansive role in future combat.

¹² For example, the US military fears that it will overload its intelligence analysts and their ability to review the information being supplied by unmanned assets if changes, to include increasing the autonomy of the systems, are not made: US Department of Defense 2012b, pp. 30–34, 82–83.

¹³ US Department of Defense 2012b, p. 1 (‘Enable humans to delegate those tasks that are more effectively done by computer... thus freeing humans to focus on more complex decision making’).

¹⁴ Sharkey 2012, p. 110 (observing that ‘armed robots are set to change the pace of battle dramatically in the coming decade. It may not be militarily advantageous to keep a human in control of targeting’).

¹⁵ For example, the US is seeking to greatly expand its use of autonomy: US Department of Defense 2012b, pp. 1–3.

13.2.2 Future Technological Possibilities

The forecast for continued improvements in autonomous capabilities seems optimistic, but predicting exactly how technology might be developed for future weapon systems is difficult. One safe assumption is that the processing power of future AWS's on-board computers will be dramatically faster and more capable than anything presently appearing on an unmanned system. As computing capabilities improve, AWS will increasingly be embedded with advanced artificial intelligence applications. These programs will likely feature a branch of artificial intelligence known as general or strong artificial intelligence, which will enable the systems to independently react to complex problems.¹⁶ Systems powered by this strong artificial intelligence will adapt and learn from their experiences and their environment.¹⁷ In fact, some contend that this feature may ultimately help AWS 'behave more ethically and far more cautiously on the battlefield than a human being'.¹⁸

Another likely facet of AWS is that they may not resemble contemporary unmanned systems. Radical changes in form and shape may be possible because of future computers, which will not only be faster and more powerful but also tremendously more compact. Thus, AWS of the future can be expected to operate both at dramatically increased ranges and without human interaction for extended periods. For instance, the US is developing designs for an anti-submarine warfare vessel capable of hunting enemy submarines autonomously for up to 3 months.¹⁹ Other AWS will be smaller, more expendable, and able to operate collaboratively as part of a swarm. In fact, swarm technology holds great promise for rapidly engaging and overwhelming an enemy.²⁰ Regardless of the exact form these technological advances yield, future AWS will possess almost unimaginable increases in capability.

¹⁶ Human-like cognitive abilities are not the equivalent of human abilities. Consensus does not exist as to if and when general artificial intelligence might become available. Computer scientist Noel Sharkey doubts that artificial intelligence advances will achieve human-like abilities in even the next 15 years: Sharkey 2011, p. 140.

¹⁷ Singer 2009, p. 74; Kellenberger 2011, p. 27.

¹⁸ Kellenberger 2011, p. 27.

¹⁹ US Defense Advanced Research Projects Agency 2013. Note, however, that at least initially the vessel is designed to require human approval before launching an attack. Similar underwater systems are being developed by the US Navy to conduct de-mining operations: Ackerman 2013.

²⁰ US Air Force 2009, p. 16 (stating that '[a]s autonomy and automation merge, [systems] will be able to swarm... creating a focused, relentless, and scaled attack'). The US Air Force's Proliferated Autonomous Weapons may represent an early prototype of future swarming systems. See, Singer 2009, p. 232; Alston 2011, p. 43.

13.3 Legal Implications of Autonomous Weapon Systems

If AWS are indeed to become a technological reality to be used in battle, it is vital to examine the relevant law that would be applied. Unquestionably, the law that would govern AWS is the law of armed conflict. The law of armed conflict has evolved over time due to changes in weaponry and tactics. Whenever a new weapon, such as an autonomous weapon, is developed and considered for fielding, two distinct aspects of the law need to be fully scrutinised: weapons law and targeting law.²¹ The former concentrates on whether the weapon itself is lawful *per se*. The latter focuses strictly on the prohibited uses of the weapon system. Both aspects must be satisfied before a weapon can be sent into battle.

13.3.1 Weapons Law Rules

Two separate weapons law rules comprise the heart of any examination of the lawfulness of a weapon system *per se*. The first rule prohibits any weapon system that is indiscriminate by its very nature. Weapons are indiscriminate by nature when they cannot be aimed at a specific objective and would be as likely to strike civilians as they would combatants. Generally considered reflective of customary international law,²² this rule is codified in Article 51(4)(b) of *Additional Protocol I*.²³ As a customary rule, all states, even those not a party to the Protocol, are bound to obey this prohibition against indiscriminate attack. The fact that the autonomous system, as opposed to a human controlled system, may make the final targeting decision is irrelevant under this rule. So long as the autonomous weapon can be supplied with sufficiently reliable and accurate data to enable it to be directed at a specific military target, the weapon system would not be indiscriminate by nature, and thus not unlawful *per se*.

The second rule, codified in Article 35(2) of *Additional Protocol I* and reflective of customary international law,²⁴ prohibits any weapon that causes unnecessary suffering or superfluous injury. The prohibition seeks to prevent inhumane or needless injuries to combatants. Warheads filled with glass are a classic example of an unlawful weapon under this rule. Such warheads unnecessarily complicate

²¹ *Legality of the Threat or Use of Nuclear Weapons (Advisory Opinion)*, ICJ Reports 1996, 226 ('*Nuclear Weapons*') p. 257 para 78; Schmitt and Thurnher 2013, p. 243. Schmitt 2013, p. 8. Targeting law is often also referred to as the rules that apply to the conduct of hostilities.

²² Henckaerts and Doswald-Beck 2005, r. 7. See also, Cadwalader 2011, p. 153.

²³ *Protocol Additional to the Geneva Conventions of 12 August 1949 relating to the Protection of Victims of International Armed Conflicts (Protocol I)*, 8 June 1977, 1125 UNTS 3 (entered into force 7 December 1978) ('*Additional Protocol I*').

²⁴ Henckaerts and Doswald-Beck 2005, r. 70; *Nuclear Weapons*, para 78; Cadwalader 2011, p. 157.

medical treatment and thus violate the prohibition. Except in the unlikely scenario where a state equips its AWS with such unlawful munitions, this rule will generally not impede the creation of AWS. While AWS can obviously only be armed with weapons and ammunition that do not cause unnecessary suffering or superfluous injury, autonomous features by themselves would not affect or violate this prohibition.

To ensure both aforementioned rules are satisfied, a state wishing to deploy a new weapon must perform a legal review. Codified in Article 36 of *Additional Protocol I*, the legal review requirement ensures that the weapon is not indiscriminate and that it would not cause unnecessary suffering or superfluous injury. The review further determines whether any other provision of the law of armed conflict might prohibit the use of the weapon. This legal review of weapons and weapon systems is generally considered a rule of customary international law. Thus, such reviews are required by all states, including those not party to the Protocol.²⁵ Additionally, if a weapon is substantially modified after fielding, then a further legal review is necessary. Clearly, prior to use, any AWS would require a legal review.

13.3.2 Targeting Law Requirements

Assuming a specific autonomous weapon satisfies the weapons law rules described above, it must still be examined under targeting law to ascertain whether the use of the weapon system might be prohibited. This analysis requires examination of three key requirements of the law of armed conflict: distinction, proportionality, and precautions in the attack. Even if a weapon satisfies the weapons law rules discussed above, it may still not be deployed if its use violates any one of these three targeting law tenets.

The first targeting law requirement is distinction. Distinction is recognised as a ‘cardinal’ principle of the law of armed conflict.²⁶ Reflective of customary international law, distinction requires a combatant to differentiate between combatants and civilians, as well as between military and civilian objects.²⁷ The rule is codified in Article 48 of *Additional Protocol I* with accompanying rules in Articles

²⁵ A legal review requirement is generally considered customary only with respect to the means of warfare, namely weapons and weapon systems. *Additional Protocol I*, Article 48 also requires a legal review of methods of warfare. An obligation to review new methods of warfare has not crystallised into customary international law: Schmitt (ed) 2013, commentary accompanying r. 48.

²⁶ *Nuclear Weapons*, paras 78–79.

²⁷ Henckaerts and Doswald-Beck 2005, r. 1; *Nuclear Weapons*, paras 78–79; Cadwalader 2011, p. 157.

51 and 52.²⁸ The principle is intended to protect civilians by directing military attacks against only military objectives,²⁹ and it unequivocally applies to AWS.

When analysing whether the use of an autonomous weapon complies with the principle of distinction, the surrounding context and environment are of critical importance. Circumstances may exist in which AWS would only need a low level ability to distinguish in order to comply with the rule. Examples include conflicts against declared hostile forces where the fighting occurs in remote areas, such as underwater, deserts, or places like the Demilitarised Zone in Korea. In less clear-cut environments, the demands on AWS to distinguish civilians from legitimate military targets are much higher. For instance, on cluttered battlefields or in urban areas, AWS may need to be equipped with robust sensor packages and advanced recognition software. Even with such cutting-edge capabilities, there could be complex situations where AWS are simply unable to fulfill this requirement and, therefore, could not lawfully be used. In the end, AWS may only be lawfully used if the systems are able to reasonably distinguish between combatants and civilians (and between military objectives and civilian objects), given the specific circumstances of the battlefield ruling at the time.

The second targeting law requirement is proportionality, which requires combatants to examine whether the expected collateral damage from an attack would be excessive in relation to the anticipated military advantage. This complex principle is reflective of customary international law,³⁰ and is codified in both Article 51(5)(b) and Article 57(2)(iii) of *Additional Protocol I*.³¹ To comply with the principle, AWS would need to be able to estimate the expected amount of collateral civilian damage that might occur as a result of an attack. Modern militaries have developed a procedure, known as the Collateral Damage Estimation Methodology, for making these required estimates.³² The methodology relies on objective and scientific criteria, and, as such, AWS should undoubtedly be able to conduct this quantitative analysis. The next step in the proportionality analysis, however, will be more complicated for AWS.

If any civilian casualties are likely to result from an attack, the proportionality rule next requires AWS to compare that amount of collateral harm against the military advantage anticipated to be gained from destroying the target. This step may present challenges for AWS, because the military advantage of a particular target is contextual. These determinations are generally made on a case-by-case

²⁸ *Additional Protocol I*, Articles 48, 51–52.

²⁹ For details, see, for example, Schmitt 2012a.

³⁰ Henckaerts and Doswald-Beck 2005, r. 14; Cadwalader 2011, pp. 157–158.

³¹ The rule specifies that an attack is indiscriminate if it is 'expected to cause incidental loss of civilian life, injury to civilians, damage to civilian objects, or a combination thereof, which would be excessive in relation to the concrete and direct military advantage anticipated': *Additional Protocol I*, Article 51(5)(b).

³² For a discussion of the methodology, see, Thurnher and Kelly 2012.

basis. It is, however, conceivable that AWS could lawfully operate upon a framework of pre-programmed values. The military operator setting these values would, in essence, pre-determine what constitutes excessive collateral damage for a particular target. Given the separation of the operator from the AWS and the fact that the judgements would be made in advance, these values would invariably need to be set at extremely conservative ends to comply with the rule.³³ Military controllers might also help AWS comport with this principle by establishing other controls, such as geographic or time limits on the use of these systems. States will have to diligently sort through these thorny proportionality issues prior to using an autonomous weapon on the battlefield.

The third and final targeting law requirement is the obligation to take feasible precautions in the attack. Customary in nature and codified in Article 57 of *Additional Protocol I*,³⁴ these precautions apply to the use of AWS and may present challenges for states wanting to deploy them. One such challenge is the requirement to do everything feasible to choose a means of attack ‘with a view to avoiding, and in any event minimizing’ collateral damage.³⁵ Feasible, in this discussion, means ‘that which is practicable or practically possible, taking into account all circumstances prevailing at the time, including humanitarian and military considerations’.³⁶ Under some circumstances, this precaution may prohibit the use of AWS if instead a different system could feasibly perform the mission and better protect civilians without sacrificing military advantage. Conversely, there may be circumstances where the use of AWS would be required, such as when their use is feasible and would offer greater protection to civilians. Another challenge may be posed by the obligation to do everything feasible to verify that a target is a military objective.³⁷ In many cases, the advanced recognition capabilities of AWS would be sufficiently precise and reliable to fulfill this requirement. Yet at other times, depending on the situation and what is practically possible, a force may have to augment AWS with other sensors to help validate the target. Of course, one must always be mindful that the standard in dealing with these rules on feasibility, as with other rules of the law of armed conflict, is reasonableness. AWS should not be held to an absolute standard or required to do more than is expected of manned or human controlled systems.

³³ Depending on future technological advances, sliding scale-type algorithms or mechanisms may be developed to allow AWS to adjust from those established baselines on their own based upon changes that the systems identify on the battlefield.

³⁴ Henckaerts and Doswald-Beck 2005, r. 15; Cadwalader, pp. 161–162.

³⁵ *Additional Protocol I*, Article 57(2)(a)(ii).

³⁶ Humanitarian Policy and Conflict Research 2009, p. 38.

³⁷ *Additional Protocol I*, Article 57(2)(a)(i).

13.4 Legal Lacunae in Autonomous Weapon Systems?

While the above section outlined the basic legal standards applicable to AWS and their use, this section delves more deeply into two areas of the law that may need further development as AWS become operational. The first topic is the subjective nature of targeting decisions and how those judgements might be made in the context of autonomous systems. The second is the issue of responsibility with respect to actions by AWS. Both subjects showcase the complex and unique legal implications of AWS.

13.4.1 Subjectivity in Targeting

Subjectivity plays a significant role in various facets of targeting law. When analysing proportionality, for instance, the military operator ordering the strike must subjectively decide the value of the target from a military advantage perspective. That person must also subjectively evaluate whether the expected collateral harm is excessive in relation to the anticipated military advantage to be gained. Similarly, with all of the required precautions in attack, there are inherently subjective judgements that must be made regarding whether all feasible precautions have indeed been taken.

It is doubtful that AWS will be able to make these subjective determinations themselves in the foreseeable future, even with the most optimistic projections for artificial intelligence advancements. Many opponents of AWS and some scholars contend that the systems are unlawful because they lack that ability to make subjective determinations.³⁸ This view is somewhat misguided, however, by failing to fully appreciate how the AWS targeting process will actually occur. In autonomous attacks, the main targeting decisions remain subjective, and those value judgements will continue to be made exclusively by humans. However, the subjective choices may be made at an earlier stage of the targeting cycle than with the more traditional human controlled systems. Sometimes these judgement calls will be made before the AWS are even launched. This difference does not necessarily make AWS unlawful. On the contrary, it merely represents a new way of looking at the subjectivity requirements.

To comply with the law, humans will need to inject themselves at various points into the process and make the necessary subjective determinations. The first such point is when a military operator programs the autonomous weapon. Depending on the weapon system's sophistication, this may transpire in the design phase, before launching the system, or perhaps even remotely during the mission. The controller must subjectively decide what numerical or other values to assign to

³⁸ Herbach 2012, pp. 17–19; Wagner 2012, pp. 121–122. See generally, Gillespie and West 2010, pp. 13–20.

targets as a guide for the autonomous system during its mission. In providing attack criteria or thresholds, the human operator is framing the environment within which the autonomous weapon will operate. The human operator is essentially providing the subjective answers in advance. Then, with that guidance embedded into its software, the autonomous system will be tasked with making objective calculations about how to perform on the battlefield.³⁹ For example, if an autonomous weapon objectively calculates that a potential strike will cause more collateral damage than the human controller has authorised, then the weapon would refrain from launching the attack and seek additional guidance or continue its mission elsewhere. In the end, it is the human, not the autonomous system, who makes the qualitative and subjective choices.

Another critical point in the process is when a military operator orders an autonomous weapon into battle. That choice is clearly a subjective one. The operator must personally decide whether the autonomous weapon can perform lawfully given the specific battlefield situation. To make such a judgement, the military controller must be thoroughly familiar with the system's particular capabilities and must know what embedded values have been pre-programmed into it. Given what the operator knows about the battlefield environment and how the autonomous weapon is programmed to react in that given environment, he or she must subjectively determine whether AWS are the correct weapons for the given mission. Operators must be certain that the AWS are expected to perform in compliance with the law of armed conflict before ordering the systems into such a situation.

These human subjective decisions will ultimately be examined for reasonableness. Timing plays a pivotal role in any measure of reasonableness with AWS. The longer the amount of time between the last human operator input and the autonomous strike itself, the greater the risk of changes on the battlefield resulting in an unanticipated action. The greater the risk, the less reasonable the decision to deploy AWS into that battle generally becomes. Certainly, the risk could be lowered if the autonomous weapon is capable of regularly submitting data about the environment back to a human operator who could potentially adjust the engagement criteria. This may not always be an option, however.⁴⁰ In the end, the human will be expected to make a reasonable decision about the appropriate amount of risk in using the autonomous systems.

The fact that AWS will not be making subjective decisions themselves should not affect the lawfulness of the systems. On the contrary, the legal requirements will be met by the subjective human input throughout the targeting process. Given

³⁹ It is important to note that the objective decisions made by AWS are distinct from the subjective ones required by the law of armed conflict. Any objective criteria are more akin to Rules of Engagement that direct the autonomous weapon's actions than to legal thresholds. These operational constraints can and would likely be set at a more stringent level than would be allowed by law.

⁴⁰ It is less likely that systems that operate underwater or in areas where communications jamming is prevalent will be able to have their subjective values adjusted during a mission.

that much of this input may occur early in the targeting cycle, however, a new way of looking at the subjective requirements of targeting law may be required.

13.4.2 Responsibility

In many ways, AWS represent a greater separation of humans from the battlefield. Therefore, significant questions arise when one looks to assess legal responsibility for battlefield conduct.⁴¹ Opponents of AWS argue that the removal of humans from the final targeting decisions prevents the proper assignment of legal responsibility.⁴² This position fails to take into account the full involvement of human commanders in the overarching targeting process. Contrary to the critics' concerns, humans can be held legally responsible for the results of AWS attacks, even when they are not controlling the specific actions of the system.⁴³

Some responsibility issues are relatively straightforward. An individual who intentionally programs AWS to engage in actions that amount to a war crime would certainly be liable. Likewise, an individual would be responsible for using a system in an unlawful manner, such as deploying AWS that are incapable of distinguishing combatants from civilians into areas where civilians are expected to be located. Superior commanders of such individuals could also be held responsible if they knew or should have known about the deliberate programming or unlawful use of the system and did not try to stop the action.⁴⁴

Other AWS responsibility issues are more complex. Humans can be held responsible for the decisions they make related to programming the weapon and to deploying it onto a battlefield in the circumstances. A human could also be held responsible for the underlying subjective targeting decisions that laid the foundation for the ultimate strike. These actions would be measured for reasonableness. As was discussed above, there are two components to the subjective targeting decisions that must be evaluated for reasonableness. First, the decision to deploy the autonomous weapon under the circumstances must be reasonable. Second, the expected length of time from the launch of the system to the strike on the target also needs to be reasonable. Given these unique challenges, a new, broader way of looking at the problem may be called for with respect to legal responsibility

⁴¹ Alston 2011, p. 51 (stating that '[o]ne of the most important issues flowing from increased automation is the question of responsibility for civilian casualties or harm or violations of the laws of war'). See also, Coughlin 2011, pp. 86–90.

⁴² Human Rights Watch 2012, p. 42.

⁴³ Fenrick 2010, p. 505.

⁴⁴ See, for example, *Geneva Convention Relative to the Protection of Civilian Persons in Time of War*, 12 August 1949, 75 UNTS 287 (entered into force 21 October 1950), Article 146; *Additional Protocol I*, Articles 86–87; *Rome Statute of the International Criminal Court*, 17 July 1998, 2187 UNTS 90 (entered into force 1 July 2002), Articles 25(3)(b) and 28. The law of armed conflict further imposes a duty to investigate possible war crimes: Schmitt 2011, pp. 31–84.

for the deployment of AWS. In the end, however, these human judgements, rather than the autonomous actions of pulling the trigger or pushing the button, are the critical ones in any issue of responsibility.

13.5 Conclusion

Before they have even been developed for use on a battlefield, AWS have sparked intense controversy. In this emerging debate, the law has often been conflated with policy, morality, and ethical concerns. This chapter sought only to address and distill out the legal issues surrounding AWS. If AWS development continues unabated, the law of armed conflict and its weapons and targeting rules will play a prominent role in shaping the systems. The unique application of autonomous targeting may influence the traditional interpretation of some of these rules, particularly in the areas of subjectivity and responsibility. As has been the case with other transformative developments in weaponry throughout history, the law will adapt and evolve as needed. In general, however, autonomy and the law of armed conflict are not inconsistent or incompatible.

In actuality, AWS and their use will likely be deemed lawful in many scenarios. When examining AWS under weapons law rules, there is little evidence to suggest that the systems would be unlawful *per se*. After all, a weapon that is autonomous is no more likely to cause unnecessary suffering or superfluous injury than any other type of weapon. Nor do the autonomous facets of a weapon prevent it from being directed at specific military targets. While targeting law requirements of distinction, proportionality, and precautions in attack do present more obstacles for the use of AWS, the challenges will not be insurmountable. There will undoubtedly be situations involving complex battlefields where AWS will be unable to comply with those principles and thus prohibited from being used. However, the frequency of such situations is likely to be diminished over time as technology improves.

Even as technology advances and autonomous systems gradually take greater control over actions, humans will continue to serve a critical function. Military operators are essential to making the subjective value judgements as required by the law. This role cannot be abdicated. While the subjective decision-making may occur earlier in the targeting process than has traditionally been the case, these judgements can nonetheless satisfy the legal requirements. Operators should further expect to be held responsible for the reasonableness of those decisions. Any human recklessness or criminal misuse of AWS will trigger the same war crimes accountability mechanisms that already exist under the law.

Autonomous weapons will profoundly affect how states fight wars. Weapons that can independently select and engage targets will prove invaluable in the frenetic pace of future warfare. Still, one should expect the moves towards fully autonomous weapons to be incremental and subtle. Systems will gradually become ever more autonomous, and humans will slowly begin to play a smaller role in the execution of actions. Ultimately, every shift towards autonomy on the battlefield

will remain dependent on whether technology can indeed deliver autonomous systems discerning and sophisticated enough to comply with the law of armed conflict.

References

- Ackerman S (2013) Navy preps to build a robot ship that blows up mines. www.wired.com/dangerroom/2013/01/robot-mine-sweeper/. Accessed 26 February 2013
- Alston P (2011) Lethal robotic technologies: the implications for human rights and international humanitarian law. *J Law Inf Sci* 21:35–60
- Cadwalader G (2011) The rules governing the conduct of hostilities in Additional Protocol I to the Geneva Conventions of 1949: a review of relevant United States references. *Yearb Int Humanit Law* 14:133–171
- Coughlin T (2011) The future of robotic weaponry and the law of armed conflict: irreconcilable differences? *Univ Coll London Jurisprudence Rev* 17:67–99
- Dahm W (2012) Killer drones are science fiction. *Wall Street J*, 15 Feb 2012, A. 11
- Fenrick W (2010) The prosecution of international crimes in relation to the conduct of military operations. In: Gill T, Fleck D (eds) *The handbook of the law of military operations*. Oxford University Press, Oxford, pp 501–514
- Gillespie T, West R (2010) Requirements for autonomous unmanned air systems set by legal issues. *Int C2 J* 4(2):1–32
- Henckaerts J, Doswald-Beck L (2005) *Customary international humanitarian law*. Cambridge University Press, Cambridge
- Heintschel von Heinegg W (2011) Concluding remarks. In: Heintschel von Heinegg W, Beruto GL (eds) *International humanitarian law and new weapon technologies*. International Institute of Humanitarian Law, Sanremo, pp 183–186
- Herbach J (2012) Into the caves of steel: precaution, cognition and robotic weapon systems under the law of armed conflict. *Amsterdam Law Forum* 4(3):3–20
- Heyns C (2013) Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions on lethal autonomous robotics. UN Doc A/HRC/23/47
- Human Rights Watch (2012) Losing humanity: the case against killer robots. www.hrw.org/sites/default/files/reports/arms1112ForUpload_0_0.pdf. Accessed 27 Feb 2013
- Humanitarian Policy and Conflict Research (2009) Manual on international law applicable to air and missile warfare. www.ihlresearch.org/amw/manual. Accessed 28 June 2013
- IEEE (2012) Look ma, no hands. www.ieee.org/about/news/2012/5september_2_2012.html. Accessed 26 Feb 2013
- Kellenberger J (2011) Keynote address. In: Heintschel von Heinegg W, Beruto GL (eds) *International humanitarian law and new weapon technologies*. International Institute of Humanitarian Law, Sanremo, pp 23–27
- Poitras C (2012) Smart robotic drones advance science. <http://today.uconn.edu/blog/2012/10/smart-robotic-drones-advance-science/>. Accessed 27 Feb 2013
- Public Broadcasting Service (2011) Smartest machines on earth. (transcript) <http://www.pbs.org/wgbh/nova/tech/smarest-machine-on-earth.html>. Accessed 27 Feb 2013
- Russell S, Norvig P (2010) *Artificial intelligence: a modern approach*, 3rd edn. Prentice Hall, Upper Saddle River
- Schmitt MN (2011) Investigating violations of international law in armed conflict. *Harv Natl Secur J* 2:31–84
- Schmitt MN (2012a) Discriminate warfare: the military necessity-humanity dialectic of international humanitarian law. In: Lovell DW, Primoratz I (eds) *Protecting civilians during*

- violent conflict: theoretical and practical issues for the 21st century. Ashgate, Farnham, pp 85–102
- Schmitt MN (2012b) Autonomous weapon systems and international humanitarian law: a reply to the critics. *Harvard National Security Journal Features*. <http://harvardnsj.org/wp-content/uploads/2013/02/Schmitt-Autonomous-Weapon-Systems-and-IHL-Final.pdf>. Accessed 27 February 2013
- Schmitt MN (ed) (2013) Tallinn manual on the international law applicable to cyber warfare. International Group of Experts at the Invitation of the NATO Cooperative Cyber Defence Centre of Excellence/Cambridge University Press, Cambridge
- Schmitt MN, Thurnher J (2013) ‘Out of the loop’: autonomous weapon systems and the law of armed conflict. *Harv Natl Secur J* 4:231–281
- Sharkey N (2011) Automating warfare: lessons learned from the drones. *J Law Inf Sci* 21:140–154
- Sharkey N (2012) Drones proliferation and protection of civilians. In: Heintschel von Heinegg W and Beruto GL (eds) *International humanitarian law and new weapon technologies*. International Institute of Humanitarian Law, Sanremo, pp 108–118
- Singer PW (2009) *Wired for war: the robotics revolution and conflict in the twenty-first century*. Penguin Press, New York
- Thurnher J, Kelly T (2012) Collateral damage estimation. US Naval War College video. www.youtube.com/watch?v=AvdXJV-N56A&list=PLam-yp5uUR1YEwLbqC0IPrP4EhWOeTf8v&index=1&feature=plpp_video. Accessed on 26 February 2013
- US Air Force (2009) Unmanned aircraft systems flight plan 2009–2047. Headquarters Department of the Air Force, Washington DC
- US Defense Advanced Research Projects Agency (2013) DARPA’s anti submarine warfare game goes live. www.darpa.mil/NewsEvents/Releases/2011/2011/04/04_DARPA's_Anti-Submarine_Warfare_game_goes_live.aspx. Accessed on 26 February 2013
- US Department of Defense (2009) FY2009–2034 unmanned systems integrated roadmap. Government Printing Office, Washington DC
- US Department of Defense (2012a) Directive 3000.09: autonomy in weapon systems. Government Printing Office, Washington DC
- US Department of Defense (2012b) Task force report: the role of autonomy in DoD systems. www.fas.org/irp/agency/dod/dsb/autonomy.pdf. Accessed 26 Feb 2012
- Wagner M (2012) Autonomy in the battlespace: independently operating weapon systems and the law of armed conflict. In: Saxon D (ed) *International humanitarian law and the changing technology of war*. Martinus Nijhoff, Leiden, pp 99–122