WHY WORDS MATTER:

THE REAL WORLD CONSEQUENCES OF DEFINING AUTONOMOUS WEAPONS SYSTEMS

*Michael C. Horowitz**

I. INTRODUCTION

A debate about the potential development and deployment of autonomous weapon systems (AWS) by countries, and potentially non-State actors, around the world first emerged in 2012. The U.S. Department of Defense (DoD) released a directive on autonomy that delineated AWS from other types of weapons,[1] and a report by Human Rights Watch and Harvard Law School's International Human Rights Clinic made the case for banning these "killer robots."[2] Since 2012, the Convention on Certain Conventional Weapons (CCW) at the United Nations (U.N.) has discussed the issue of autonomous weapons twice, with a third dialogue planned for April 2016.

The debate right now is framed by fear that AWS could represent a new, dangerous class of weapons fundamentally distinct from the weapons of today. Thus, groups such as the Campaign to Stop Killer Robots, a collaboration between dozens of non-government organizations (NGOs), argue that the international community should ban these weapons before they are deployed.[3] The problem is that no one seems quite sure what an autonomous weapon *is*, and thus what the international community should be discussing.

Decisions by the international community about how to define an autonomous weapon will have important consequences. These decisions will determine whether the international community views autonomous weapons solely as weapons of the future, versus including weapons in the present, and thus what is under discussion in debates about potential regulations or bans on autonomous weapons. This, in turn, has important consequences for whether a ban could be successfully negotiated and enforced.[4]

---

*Associate Professor of Political Science and Associate Director of Perry World House, University of Pennsylvania, February 2016.

1. U.S. DEP'T OF DEF., DIR. 3000.09, AUTONOMY IN WEAPONS SYSTEMS (NOV. 21, 2012) [hereinafter DOD DIRECTIVE 3000.09].

2. INT'L HUMAN RIGHTS CLINIC, LOSING HUMANITY: THE CASE AGAINST KILLER ROBOTS (2012), https://www.hrw.org/sites/default/files/reports/arms1112_ForUpload.pdf [hereinafter LOSING HUMANITY: THE CASE AGAINST KILLER ROBOTS].

3. *About Us*, CAMPAIGN TO STOP KILLER ROBOTS, https://www.stopkillerrobots.org/about-us/ (last visited Apr. 12, 2016).

4. Rebecca Crootof, "The Killer Robots Are Here: Legal and Policy Implications," CARDOZO L. REV. 36 (2015): 1837–1915. She argues that certain AWS characteristics make it unlikely that a ban will be an effective means of regulation.

This paper lays out three competing conceptions of what constitutes an AWS, and proposes a new, function-based way to distinguish between types of autonomous weapons: separating them into munitions, platforms, and operational systems. More practical distinctions between types of autonomous weapons can help clarify what is new, and what is not, in discussions of autonomy in military systems. The article concludes by arguing that this distinction, despite the inherent limitations of any definitional exercise, can improve the ability of the international community to discuss the ethical, strategic, and legal challenges associated with autonomous weapons.[5]

II. WHAT IS AN AUTONOMOUS WEAPON?

A challenge in the attempt to understand the ethical, legal, moral, strategic, and other issues associated with AWS is a basic lack of agreement over what an autonomous weapon actually is. The most prominent definition of an AWS comes from the DoD, which defines an AWS as "[a] weapon system that, once activated, can select and engage targets without further intervention by a human operator."[6] An AWS is distinguished from a semi-autonomous weapon, which is "[a] weapon system that, once activated, is intended to only engage individual targets or specific target groups that have been selected by a human operator."[7] The Campaign to Stop Killer Robots and groups such as Human Rights Watch have used similar definitions, though some NGOs now define autonomous weapons as those that lack meaningful human control.[8] Rebecca Crootof, Resident Fellow at the Information Society Project at Yale University, adds further detail by specifying that an AWS is "a weapon system that, based on conclusions derived from gathered information and preprogrammed constraints, is capable of independently selecting and engaging targets."[9] This definition adds the processing of information to the definition to describe what the machine is doing, along with the word "independent" (in the attempt to distinguish AWS from "automatic" weapons, like landmines, that are merely triggered).[10] I have defined AWS as "weapon systems that, once activated, are designed to select and engage targets not previously designated by a human."[11] This adds "not previously designated by a

5. There are inherent limitations because definitions almost inevitably lead to discussions of cases at the margins that may fall between the cracks of a definition.

6. DOD DIRECTIVE 3000.09, *supra* note 1, at 13.

7. *Id*. at 14.

8. *See, e.g*., LOSING HUMANITY: THE CASE AGAINST KILLER ROBOTS, *supra* note 2, at 2–3. Rebecca Crootof, "War, Responsibility, and Killer Robots,"40 N. CAROLINA J. OF INT'L L. AND COM. REG. 4 (2015). Note that this also raises the question of what meaningful human control means, which is briefly addressed below but also an important topic for future research. The U.S. government has used the related phrase "appropriate human judgment." 2016 Meeting on Experts on LAWS, U.S. Delegation Opening Statement (Apr. 11, 2016), http://www.unog.ch/80256EDD006B8954/(httpAssets)/EFF7036380934E5EC1257F920057989A/$file/2016_LAWS+MX_GeneralExchange_Statements_United+States.pdf.

9. Crootof, *supra* note 4.

10. *Id*. at 1857.

11. MICHAEL C. HOROWITZ, THE ETHICS AND MORALITY OF ROBOTIC WARFARE: ASSESSING THE DEBATE OVER AUTONOMOUS WEAPONS 5 (2016), http://www.michael

human" to account for when humans lawfully designate groups of targets for attack without knowing precisely which munition will strike precisely which target.[12] It is designed to account for the fact that lawful warfare can traditionally involve salvos and humans do not always have direct control over a munition from point of launch to point of impact, though it does not resolve uncertainty on what it means to select and engage.[13]

Beyond this, there are a number of distinctions that various groups use to delineate within the potential category of AWS, though they are not necessarily fundamental to the definition. For example, the DoD also defines human-supervised autonomous weapons (e.g. autonomous weapons where a human has an override switch).[14] The Future of Life Institute, an outreach organization that works to ensure that technology is helpful rather than harmful to humanity, wrote an open letter calling for a ban on AWS and focuses on "a ban on offensive autonomous weapons beyond meaningful human control,"[15] adding "offensive" and "beyond meaningful human control" as modifiers to express their concern. What meaningful human control means also requires clarification, however. How much information and direct supervision makes human control meaningful?

The U.N. discussions on AWS as part of the CCW focused on lethal autonomous weapon systems (LAWS), which adds the term lethal to modify "autonomous weapon system."[16] In that case, AWS that do not directly target people or military platforms carrying people might not fall under the definition. A 2015 report by Human Rights Watch and Harvard Law School's International Human Rights Clinic focused on fully autonomous weapon systems, adding the term "fully" as a modifier to presumably delineate between semi-autonomous, human-supervised autonomous, and fully autonomous weapons systems, though that is not specified.[17]

The variety of modifiers introduces complexity into discussions of AWS because different commentaries or reports might focus on slightly different technologies, depending on the specific modifiers being used. Many of these definitions, however, gloss over a fundamental challenge at the heart of defining an

---

chorowitz.com/Documents/HorowitzLAWSEthicsDraftFeb2016.pdf.

12. *Id*. at 5.

13. *Id*.

14. DOD DIRECTIVE 3000.09, *supra* note 1, at 14.

15. *Autonomous Weapons: An Open Letter from AI & Robotics Researchers*, FUTURE OF LIFE INST. (July 28, 2015), http://futureoflife.org/AI/open_letter_autonomous_weapons.

16. For example, the official Convention on Certain Conventional Weapons webpage on its 2016 discussions describes it as a "[m]eeting of [e]xperts on LAWS," rather than a meeting of experts on AWS. *See* 2016 Meeting of Experts on LAWS, *supra* note 8 (stating that the meeting to be convened in 2016 is to discuss questions in the area of *lethal* autonomous weapons systems).

17. HUMAN RIGHTS WATCH, MIND THE GAP: THE LACK OF ACCOUNTABILITY FOR KILLER ROBOTS (2015), https://www.hrw.org/report/2015/04/09/mind-gap/lack-accountability-killer-robots.

autonomous weapon system—what it means to select and engage a target.[18] It is the definition of those terms that is crucial.[19]

The DoD Dictionary of Military Terms does not define the word "select."[20] According to the Oxford English Dictionary, select simply means "to choose or pick out."[21] "Engage" is a military term of art defined by the DoD as "to bring the enemy under fire."[22] While from a common sense perspective, choosing/picking what to bring under fire might seem clear, in reality it could mean several things, depending on the level of aggregation in a military engagement where one believes that selection and engagement occurs.

Take the example of a human operator considering launching missiles into an area where there are enemy naval surface ships. If that human operator launches the missiles, believing there are lawful enemy combatants in that area, one way to interpret select and engage is that those missiles, by definition are not autonomous weapons, because the human operator has decided to launch missiles into an area where the operator believes enemy combatants exist. Another interpretation, though, might view those missiles as AWS if the human operator is not individually picking out the targets, but the missile is (via a guidance system that has data about the targets programmed in and uses radar or other means to track).

Now imagine that instead of a human operator launching the missiles, a human operator believes there are enemy naval surface ships 1000 miles away, but is not sure if the ships are there. The human operator launches a robotic system programmed by algorithm to launch missiles at adversary ships if those ships are detected. The robotic system detects the ships and launches the missiles. One way to interpret select and engage for this system is that this is not, in fact, an AWS, because a human operator believed there were ships around that location and launched a weapons platform to attack those ships. The human is still doing the selection and engagement despite the uncertainty involved about the location of the ships. On the other hand, the weapon platform is run by algorithm and responsible for the details of the engagement—finding, fixing, tracking, targeting, and terminally engaging the enemy ships. How is that not an autonomous weapon?

Given this uncertainty, several potential approaches have emerged to understand what an AWS is.[23] Resolving this debate is important, because the

18. *See, e.g.*, *id*. Crootof, however, disagrees with this presumption.

19. For exceptions, see Crootof as well as Kenneth Anderson, Daniel Reisner, and Matthew C. Waxman, *Adapting the Law of Armed Conflict to Autonomous Weapon Systems*, 90 INT'L L. STUD. 386, 411 (2014).

20. U.S. DEP'T OF DEF., DICTIONARY OF MILITARY TERMS (2010) [hereinafter DICTIONARY OF MILITARY TERMS].

21. *Select*, OXFORD ENGLISH DICTIONARY (2016), http://www.oed.com/search?searchType=dictionary&q=select&_searchBtn=Search.

22. DICTIONARY OF MILITARY TERMS, *supra* note 20, at 79.

23. *See, e.g.*, Paul Scharre & Michael C. Horowitz, *An Introduction to Autonomy in Weapon Systems*, 8 (CENTER FOR A NEW AM. SECURITY, Working Paper No. 021015, 2015), http://www.cnas.org/sites/default/files/publicationspdf/Ethical%20Autonomy%20Working%20Paper_021015_v02. pdf; *see also Q&A on Fully Autonomous Weapons*, HUM. RTS. WATCH (Oct. 21, 2013), https://www.hrw.org/news/2013/10/21/qa-fully-autonomous-weapons.

ethical, legal, moral, and strategic issues generated by autonomous weapons may vary a great deal depending on how exactly one draws the line. Moreover, the potential effectiveness of any regulatory efforts could vary depending on just how many different weapon systems are considered "autonomous." In what follows, I describe a few potential approaches, and then a new way to think about what constitutes an AWS that relies on distinguishing between the specific actions the system is taking.

A. Narrow Construction

Most writing on AWS tends to start from the assumption that autonomous weapons are distinct from the weapons of today, excluding rare exceptions. As Crootof writes, "[t]here is a nearly universal consensus, among both ban advocates and skeptics, that autonomous weapon systems do not yet exist."[24] One possible exception is the close-in weapon systems such as the Goalkeeper, AK-630, and Phalanx operated by militaries around the world.[25] These systems operate guns or short-range defensive missiles that protect ships and military bases from attack.[26] They often have automatic modes so that if the number of incoming threats exceeds the ability of a human operator to target and fire against them, an operator can activate an autonomous mode to independently—and more quickly—respond to the incoming threats.[27] Over thirty countries operate systems with these characteristics.[28] These systems fall into the category of human-supervised autonomous weapons, since there is still a human that can override the system. If one believes that autonomous weapons do not exist today, that would likely exclude these systems, meaning AWS have to involve a greater degree of autonomy—perhaps meaning no human supervision.

The United Kingdom (U.K.) Ministry of Defence has defined an AWS in a way that excludes the weapons of today, thus drawing the line beyond today's weapons.[29] The U.K. Ministry of Defence then sets the bar higher, defining a truly

---

24. Crootof, *supra* note 4, at 1863.

25. *See* Scharre & Horowitz, *supra* note 23, at 8 (listing several close-in systems, where they are manufactured, and which nations operate them).

26. *Close-in Weapon System Radar*, INTEGRATED PUB., http://www.tpub.com/fcv2/19.htm (last visited Apr. 23, 2016).

27. *See Netherlands Ministry of Defence Awards Thales Goalkeeper Modification Contract*, THALES (Nov. 29, 2012), https://www.thalesgroup.com/en/content/netherlands-ministry-defence-awards-thales-goalkeeper-modification-contract (stating that the Goalkeeper is autonomous and can perform kill assessment for several targets simultaneously).

28. Scharre & Horowitz, *supra* note 23, at 21.

29. It does state at one point that the Counter-Rocket, Artillery, and Mortar (C-RAM) system is an autonomous weapon system under a conventional understanding today (C-RAM is a land-based version of the Mk-15 Phalanx), but the other uses of the term autonomous in the document to refer to a much greater level of decision-making by the machine. *Joint Doctrine Note 2/11: The UK Approach to Unmanned Aircraft Systems*, U.K. MINISTRY OF DEF., at 5-2 (Mar. 30, 2011), https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/33711/20110505JDN_211_UAS_v2U.pdf.

autonomous system as a system that "will, in effect, be self-aware."[30] It states that:

> [a]n autonomous system is capable of understanding higher level intent and direction. From this understanding and its perception of its environment, such a system is able to take appropriate action to bring about a desired state. It is capable of deciding a course of action, from a number of alternatives, without depending on human oversight and control, although these may still be present.[31]

The benefit of drawing the line far beyond current weapon systems, essentially at the point of advanced artificial intelligence, is that it clearly accomplishes the goal of distinguishing AWS from today's weapons. A downside, as Crootof points out, is that it sets the bar so far beyond the technological realm of the possible that the category becomes less useful from a regulatory perspective.[32] Moreover, a narrow definitional approach may require rethinking the basic terminology associated with AWS, since it does not comport with the common sense reading of the terms "select" and "engage." Thus, if one wanted to regulate systems such as the Harpy[33] or Goalkeeper,[34] it would require a separate definitional construct.

B. Broad Construction

An alternative approach is to set aside the various modifiers above and focus on the key facet of most current definitions of AWS—the selection and engagement of targets.[35] If this definition is used and one thinks of any activities in the selection and engagement process as selection and engagement writ large, AWS likely exist today in at least two forms. First, the close-in weapon systems described above, when they have an automatic mode where the system can select

---

30. *Id*. at 2-3.

31. *Id*.

32. Crootof, *supra* note 4, at 1853; *see also* Noel Sharkey, *Automating Warfare: Lessons Learned from the Drones*, 21 J. OF L, INFO, & SCI. 140, 140–55, (2011) ("[A]part from metaphorical use, no system is capable of 'understanding' never mind 'understanding higher level intent'. This would mean that autonomous robots may not be possible in the foreseeable future. This is not just pickiness about language on my part. Correctly defining what is meant by 'autonomous' has very important consequences for the way that the military, policy makers and manufacturers think about the development of military robots.").

33. Paul Scharre, *Autonomy, "Killer Robots," and Human Control in the Use of Force – Part I*, JUST SECURITY (July 9, 2014) ("Once launched, the Harpy flies over a wide area searching for enemy radars. If it detects any radars that meet its criteria, the Harpy then dive-bombs into the radar, destroying it . . . . The person launching the Harpy does not known [sic] which particular radars are to be engaged, only that radars that meet the Harpy's programmed parameters will be engaged.").

34. *See* Crootof, *supra* note 4, at 1870 (explaining that the Goalkeeper is an automated weapons system from the Netherlands that is used to identify and eliminate incoming ballistic threats with modes that allow the system to use defensive force independently).

35. *See id*. at 1857 (explaining that a key feature of an AWS is its ability to gather and process information to identify and select targets irrespective of whether this is done through programming or with no restrictions).

and engage targets, are certainly autonomous weapons under a broader definition.[36] Even if there are ways to use them in non-autonomous fashion, that they have the capacity, as programmed, to operate autonomously means they are AWS.[37]

Second, there are some categories of existing munitions, such as Israel's Harpy, that may constitute an AWS.[38] The Israeli Harpy is a loitering cruise missile that searches for a particular radar signal from an air defense system.[39] When it detects that signal, it hones in on that system and explodes.[40] This arguably makes it an autonomous weapon, since the missile is selecting the target based on the radar signal, and then engaging the target.[41] The Brimstone missile deployed by the U.K. may also be able to operate as an autonomous weapon, though that is less clear.[42] Its seeker searches for targets based on a programmed target signature, and will only hit targets identified in its system memory as valid and within a given limited geographic area.[43] While this does not meet the U.K. definition of autonomy described above, since it is not self-aware or evaluating intent, a broad definition of an autonomous weapon could include the Brimstone and other homing munitions.[44]

---

36. *See id*. at 1858–59 (asserting that although the U.S. Phalanx Close-In Weapons Systems (CIWS) includes a manual override option, it contains elaborate programming that processes information and selects targets in real time, and further noting that the CIWS also has a "casualty" mode which allows the CIWS to function independently of operators).

37. *See id*. (asserting that AWS that have manual overrides have modes which allow the system to operate independently)*; see also* Franz-Stefan Gady, *Meet Russia's New Killer Robot*, THE DIPLOMAT (July 21 2015), http://thediplomat.com/2015/2007/meet-russias-new-killer-robot/ ("The robot's targeting mechanism works automatically without human assistance . . . .").

38. *See* Crootof, *supra* note 4, at 1871 (asserting that the Israeli Harpy can identify and attack a target without direct engagement by an operator)*; see also* Peter J. Spielmann, *Israel Killer Robots Could Be Banned Under UN Proposal*, THE TIMES OF ISRAEL (May 3, 2013), http://www.timesofisrael.com/israeli-killer-robots-could-be-banned-under-un-proposal/ ("Israel's Harpy[is] a 'Fire-and-Forget' autonomous weapon system designed to detect, attack and destroy radar emitters.").

39. *See* Crootof, *supra* note 7, at 1871 ("The Israeli Harpy Loitering Weapon 'detects, attacks and destroys enemy radar emitters.' Unlike a semi-autonomous weapon, '[t]he person launching the Harpy does not know . . . which particular radars are to be engaged, only that radars that meet the Harpy's programmed parameters will be engaged.'").

40. *See id*.

41. *See id*. at 1878 ("Certain fire-and-forget or lock-on-after-launch weapon systems like the Israeli Harpy or U.K. Brimstone are employed . . . [by] autonomously selecting and engaging targets.").

42. *See id*. at 1870–71 (discussing how the U.K. Brimstone may be considered an AWS under a broader definition that includes weapons that can automatically identify and engage targets).

43. *Id*. at 1871 ("When in search mode, 'Brimstone's mmW seeker searches for targets in its path, comparing them to a known target signature in its memory. The missile automatically rejects returns which do not match . . . and continues searching and comparing until it identifies a valid target.'").

44. *See id*. ("It is unclear whether the operator programs the Brimstone to seek out a specific radar target or ones with certain characteristics; to the extent the latter is true, the Brimstone would be an autonomous weapon system".).

A broad definitional construct makes clearer what selection and engagement of targets means than the narrow construct, because it does not require conceptualizing what it means for a machine to become self-aware. A practical challenge is that this definition runs counter to arguments made by both the Campaign to Stop Killer Robots and governments such as the United States and the United Kingdom, which argue that autonomous weapon systems generally do not exist today.[45] Militaries are unlikely to get rid of weapon systems they have used without controversy for a long time, and NGO campaigns seem to recognize this issue.[46] Moreover, it would complicate attempts to regulate said weapons if one desired regulation beyond simply applying the same test used for any other weapons—whether they can comply with basic law of war requirements for necessity, distinction, and proportionality.[47] The broad definitional construct, by potentially incorporating many current weapons, also distracts from conversations about the essence of the concerns raised about AWS, which are about emerging capabilities.

C. It is The Usage, Not The Weapon

A third way to think about defining autonomous weapons is that some weapons are not inherently autonomous or not autonomous, but it depends on how they are used. Consider the case of the Advanced Medium Range Air-to-Air Missile (AMRAAM). The AMRAAM is a "fire and forget" munition used by the U.S. military and over thirty other militaries around the world.[48] Pilots use long-range radar data, not visual evidence, to determine whether a threat exists, then launch one or more missiles.[49] The AMRAAM has internal navigation and radar to find and destroy a target.[50] The AMRAAM is generally considered a semi-autonomous weapon, rather than an autonomous weapon, because a human operator believes there is a legitimate target due to long-range radar, and then

45. *See The Problem*, CAMPAIGN TO STOP KILLER ROBOTS, https://www.stopkillerrobots.org/the-problem/ (last visited Apr. 22, 2016) (arguing that AWS should be banned in the future). *See generally* LOSING HUMANITY: THE CASE AGAINST KILLER ROBOTS, *supra* note 2.

46. *See* HUMAN RIGHTS WATCH, MIND THE GAP: THE LACK OF ACCOUNTABILITY FOR KILLER ROBOTS (2015), https://www.hrw.org/report/2015/04/09/mind-gap/lack-accountability-killer-robots (discussing the roadblocks to regulation).

47. This is an even simpler approach, arguably, because it makes the definition of an autonomous weapon less important. It is the preferred solution of Anderson, Reisner, and Waxman. *See, e.g.*, Kenneth Anderson & Matthew C. Waxman, *Law and Ethics for Autonomous Weapon Systems: Why a Ban Won't Work and How the Laws of War Can*, THE HOOVER INST. (Apr. 10, 2013), http://media.hoover.org/sites/default/files/documents/Anderson-Waxman_ Law AndEthics_r2_FINAL.pdf; Anderson et al., *supra* note 19, at 411.

48. *See Advanced Medium-Range Air-to-Air Missile*, RAYTHEON, http://www.raytheon.com/capabilities/products/amraam/ (last visited Apr. 23, 2016) (discussing the versatility of AMRAAM for air-to-air and air-to-surface targetting).

49. *See* Scharre & Horowitz, *supra* note 23, at 11 (asserting that in air-to-air engagements, there is a high degree of automation, such as the fact that potential targets are identified by a computer and not by the pilot).

50. *See Advanced Medium-Range Air-to-Air Missile*, *supra* note 48 (explaining that the AMRAAM has an active guidance system and its seeker design provides the ability to find targets).

launches the missile.[51]

In general, used with the rules of engagement generally designated by responsible militaries, semi-autonomous munitions are not autonomous weapons.[52] A pilot believes there are lawful targets over the horizon and launches AMRAAM missiles in response. Even though the pilot is not specifically targeting each missile to each target, this represents appropriate human control of the sort required by the law of war. However, imagine a case where a pilot is not sure there are lawful targets over the horizon, but launches several AMRAAM missiles anyways. If there are planes flying in the general vicinity, the AMRAAM missiles lock onto those planes and destroy them. Or a case where the targets are lawful, but the pilot does not know which missile will hit which target. Depending on when exactly the seekers for each missile turn on, each plane may be attacked, or multiple missiles could hone in on the same plane. In this case, the missile is making "choices" that arguably make it an autonomous weapon.[53] According to this perspective, it is possible for many types of existing munitions to be used autonomously by an irresponsible human actor, e.g. launched into a particular geographic space without regard for whether the people in that space are all lawful combatants. Thus, it is not simply that munitions are either autonomous or not autonomous.

A benefit of this approach is that it focuses on regulating the use of force, rather than the weapon itself, keeping the legal emphasis on the law of war.[54] A potential downside is that some weapons might be more or less inherently autonomous. A hunter-killer drone that makes choices, itself, about who to target, within broad mission parameters, would clearly represent an autonomous weapon. An arrow fired by a medieval archer is clearly not an autonomous weapon. Even though the archer does not have control over the arrow after firing it, and wind could re-direct the arrow to another target, the arrow itself is not actively selecting or engaging a target.[55]

Thus, while there are some systems, in particular semi-autonomous munitions, that have the potential to be used in ways that are autonomous if a human operator does not deploy them with care, this perspective probably does not capture the vast majority of weapons. It does point out, however, a potential issue with a binary definition of an autonomous weapon system when it comes to what

---

51. Scharre & Horowitz, *supra* note 23, at 11.

52. *See, e.g.*, *Appendix E: Rules of Engagement for U.S. Military Forces in Iraq*, HUM. RTS. WATCH, https://www.hrw.org/reports/2003/usa1203/11.htm (last visited Apr. 11, 2016).

53. Or, at least, this raises questions of what constitutes autonomous target selection.

54. *See generally* Anderson & Waxman, *supra* note 47; Anderson et al., *supra* note 19.

55. *See* DOD DIRECTIVE 3000.09, *supra* note 1, at 13 ("[An AWS is a] weapon system that, once activated, can select and engage targets without further intervention by a human operator.") A medieval bow and arrow lacks the characteristics noted in this definition. *See* Michael C. Horowitz & Paul Scharre, *Meaningful Human Control in Weapon Systems: A Primer*, 5 (CENTER FOR A NEW AM. SECURITY, Working Paper No. 031315, 2015), http://www.cnas.org/sites/default/files/publications-pdf/Ethical_Autonomy_Working_Paper_031315.pdf (discussing the use of arrows which require skill to master and use).

critics might want to regulate; if the goal is to prohibit every system that could be used in an autonomous fashion, that would include a great number of precision guided munitions used by militaries today to reduce civilian casualties and more accurately target in wartime.[56]

III. DEFINING AUTONOMOUS WEAPONS: A SYSTEMS-BASED APPROACH

An alternative way to think about grappling with the issue of AWS is to shift from trying to figure out the most precise AWS definition to trying to understand what different types of AWS might do, and thus the unique challenges those AWS might raise for the lawful use of force according to the law of war, along with potential ethical and moral concerns. Military technologies that are part of the kill chain in war generally fall into one of three categories: munitions, platforms, and operational systems.[57] In what follows, I briefly describe each of these categories, and the potential implications for thinking about autonomous weapons.

A munition is the physical weapon designed to destroy a given target. For example, the DoD Dictionary of Military and Associated Terms defines a "precision guided munition" as "[a] guided weapon intended to destroy a point target and minimize collateral damage."[58] A munition is also not returnable (like an arrow or a bullet fired from a gun), meaning it is inherently a one-way weapon.[59] AWS in the munitions category arguably already exist, such as the Harpy and Brimstone.[60]

What kinds of challenges do these systems raise when it comes to crucial questions such as discrimination, law of war, and accountability? Potentially not very many, relative to semi-autonomous munitions or even "dumb" munitions. The launch of an AWS munition is done by a person who can be held accountable for the use of that weapon the same way a person can be held accountable for the use

56. Paul Scharre, *Robotics on the Battlefield Part II: The Coming Swarm*, CENTER FOR A NEW AM. SECURITY 10 (Oct. 15, 2014), http://www.cnas.org/sites/default/files/publications-pdf/CNAS_ TheComingSwarm_Scharre.pdf.

57. *See, e.g.*, DEPARTMENT OF DEFENSE DICTIONARY OF MILITARY AND ASSOCIATED TERMS (Barry Leonard ed., 2011) ("[Munition refers to a] complete device charged with explosives, propellants, pyrotechnics, initiating composition, or nuclear, biological, or chemical material for use in military operations, including demolitions. Certain suitably modified munitions can be used for training, ceremonial, or nonoperational purposes. Also called ammunition."); *Platform Information Technology Definitions for the Department of the Navy*, DEP'T OF THE NAVY 1 (Nov. 7, 2007), http://www.doncio.navy.mil/uploads/Enclosure1 _PlatformITDefinitionsforDON[2].pdf ("[Platform refers to a] vehicle, structure or person that performs a mission in support of US National Security policy; and aboard or in which a DoD national security system may be installed to support assigned missions."); *Land Combat Systems*, FED'N OF AM. SCIENTISTS, http://fas.org/man/dod-101/sys/land/ais.htm (last visited Apr. 22, 2016) (describing various automated information operational and battle systems). *See generally* Horowitz & Scharre, *supra* note 19.

58. DICTIONARY OF MILITARY TERMS, *supra* note 20.

59. *See, e.g.*, *Smaller and Smarter: Breakthroughs Point Way to Miniaturized Weapons that Never Miss*, RAYTHEON (Apr. 21, 2015), http://www.raytheon.com/news/feature/guided_ weapons.html (describing the use of these one-way munitions).

60. Crootof, *supra* note 4, at 1870–71.

of semi-autonomous "fire and forget" munitions today.[61] If the person launching the missile does not launch it in a case that meets law of war criteria concerning necessity, distinction, proportionality, etc., that person could be held accountable.[62] There is also no reason to expect that an AWS munition would be inherently indiscriminate. If it did not function in a way that could discriminate, it would be illegal in a way similar to any other non-AWS munition. For example, imagine an Israeli Harpy that honed in on a radar on top of a school, rather than a radar on top of a military base. Unless the school was a lawful target given the particular military contingency, the person who launched the Harpy in an area where it could hone on to a radar on top of a school could be held responsible. Thus, for AWS munitions, existing law of war regulations likely are sufficient, potentially with little change, to ensure weapons are used in ways that comport with the law of war and allow for accountability.[63]

Platforms are systems that launch munitions and that are inherently returnable, though they could in theory be sent on one-way missions.[64] For example, an airplane launched in a dangerous environment in which it is likely to be shot down, or even where it will accomplish its mission by flying directly into a target (e.g., kamikaze bombers in World War II)[65] is still a platform, rather than a munition, because it has the ability to return home.[66] Tanks and submarines are other examples of platforms. The only AWS platforms that seem to exist now are the close-in-weapon systems described above, though their applicability depends on the definition.

AWS platforms would raise a larger set of potential challenges than munitions, though many of them might have more to do with military effectiveness

---

61. This is not to say that accountability works perfectly today. That no person was punished for the Patriot missile system malfunction in 2003 demonstrates that accountability systems are difficult to implement in all cases. *See* Charles Piller, *Vaunted Patriot Missile Has a 'Friendly Fire' Failing*, L.A. TIMES (Apr. 21, 2003), http://articles.latimes.com/2003/apr/21/news/war-patriot21 (highlighting that this would not be a unique issue for AWS munitions).

62. *See The Law of Armed Conflict*, INT'L COMM. OF THE RED CROSS (June 2002), https://www.icrc.org/eng/assets/files/other/law1_final.pdf (describing the law of war criteria).

63. Note that this also suggests that, to the extent that a lack of human accountability is essential to the definition of an autonomous weapon, the munitions described above would not constitute autonomous weapons. *See* Scharre & Horowitz, *supra* note 23, at 8 (discussing the issue of human accountability in AWS). On related points, see Michael N. Schmitt, *Autonomous Weapon Systems and International Humanitarian Law: A Reply to the Critics*, HARV. NAT'L SECURITY J. (2013).

64. *See* DEP'T OF DEF., DEFENSE SCIENCE BOARD TASK FORCE ON MUNITION SYSTEMS RELIABILITY 14 (2005), http://www.acq.osd.mil/dsb/reports/ADA441959.pdf (discussing platforms reliability).

65. Japanese Kamikaze pilots would crash their aircrafts into Allied Ships during World War II in hopes of sinking or seriously damaging a vessel. *Remembering Japan's Kamikaze Pilots*, BBC NEWS (Feb. 26, 2014), http://www.bbc.com/news/magazine-26256048.

66. In theory, one could imagine a potential hybrid between a munition and a platform—a system that can return to base if it does not find a target, but that is the munition itself, e.g., it does not launch munitions. *See* DEP'T OF DEF., *supra* note 64, at 6 (discussing munitions and platforms systems failures and possible improvements).

than legal or ethical issues. For example, the immense programming challenge involved in having AWS platforms searching for insurgents or conducting tasks in an urban environment could raise the issue of fragility and system failures.[67] However, that is also why actors are unlikely to deploy AWS in those situations.

Establishing fair accountability systems for AWS platforms would be complicated. It would not be hard to create rules of course—the issue would be creating rules that fairly assign responsibility. If an AWS platform malfunctions, for example, is the programmer, soldier who launches the platform, or commander that launches the mission truly responsible? It might depend on the situation. One thing that can help mitigate these concerns would be clear rules and training about when deploying AWS might lead to greater operational risk—and when it would not. For example, AWS platforms deployed in clear war zones and with limited geographic scope and programming might be easier to create accountability systems for than AWS platforms with global reach and scope across zones of war and peace.

The final AWS level is that of the operational system.[68] This would involve AWS actually planning military operations in a way that completely replaces the human staff systems that plan military operations today. This is the AWS activity most akin to science fiction, and also one that militaries are least likely to pursue—since they want to maintain control over the use of force.

AWS operational systems would raise the largest degree of potential challenges. As the decision-making distance between the person activating the system and the system selecting and engaging targets increases, it makes fair accountability harder and makes a responsibility gap more likely. For example, a person would still have to activate an operational planning AWS, meaning one could establish accountability that way. However, what that AWS would do might be highly uncertain, disconnecting the person from the actual selection/engagement of targets in a way even more distant from the platform level. This is the situation where meaningful human control/appropriate human judgment would be most at risk.

Thus, across the three different potential levels of AWS, there are potentially significant differences in how complicated it would be to integrate them into existing legal and ethical frameworks for the use of force. Rather than only thinking about what distinguishes AWS from non-AWS, it therefore makes sense to think about the functions that different AWS might serve.

One benefit of this approach is that it is flexible enough to work regardless of the definitional choice made above, though it might also require three definitions – one for each level of autonomous weapon. While focusing on the function of particular potential AWS does not solve the issue of how to define them at the

67. *See* Michael C. Horowitz & Paul Scharre, *The Morality of Robotic War*, N.Y. TIMES (May 26, 2015), http://www.nytimes.com/2015/05/27/opinion/the-morality-of-robotic-war.html (arguing that humans still need some control over AWS so that they can be held accountable for AWS' actions).

68. *See* PAUL SCHARRE, AUTONOMOUS WEAPONS AND OPERATIONAL RISK (2016) (categorizing AWS according to the level of human involvement).

margins, it does provide a way to focus the discussion, potentially helping reveal which definitional approach best captures the relevant issues. Another benefit is that it makes clear why some types of AWS might be more or less problematic than others. For example, an attempt to ban all weapons with autonomous target selection defined broadly could end up prohibiting current and next generation precision munitions that make the use of force more accurate and less likely to kill civilians. Having gone through this exercise, it then may make sense to add back in modifiers like human-supervised to describe what is preferred—or not—by different NGOs or governments.

IV. CONCLUSION

Autonomous weapon systems raise novel challenges for thinking about the use of force in some ways, but not in others. One of the largest issues involved in the ongoing debate involves how exactly to define them—a critical step before beginning to talk seriously about a regulatory regime, to say nothing of the ban advocated by the Campaign to Stop Killer Robots.[69] This article presents three different ways to think about defining AWS: a broad, narrow, and usage-based approach, discussing some of the pros and cons of each. Each of these definitional approaches has upsides and downsides. A broad definition of AWS, while potentially accurate in capturing all instances in which a machine selects and engages targets, might wrap in so many current weapon systems that the category of AWS no longer describes a useful category for regulatory purposes. A narrow definition of self-aware machines making choices, while certainly ruling out the weapons of today, introduces new challenges in precisely defining terminology such that countries can predictably know where the line is between a semi-autonomous and autonomous weapon. A usage-based approach also has limitations.

While not resolving this definitional dispute, a functional approach to AWS based on the type of system could represent one way to move forward: by identifying the specific things AWS might do that could be problematic, or not. A function-based approach to AWS might also more coherently allow for the discussion to layer back in the "modifiers" used when discussing autonomous weapons, such as lethal (e.g. targeting humans or equipment with humans in it), human-supervised, fully, and others. By focusing on what it is that autonomous weapons may be doing on the battlefield and how that compares to the weapons of today, scholars and policymakers can get a better sense of the appropriate way to think about the AWS issue. And any efforts to move forward must begin with a firm understanding of how advanced weapons operate today. For example, whether to simply apply the law of war to uses of force, not systems,[70] develop "war torts,"

69. *See The Problem*, CAMPAIGN TO STOP KILLER ROBOTS, https://www.stopkillerrobots.org/the-problem/ (last visited Apr. 23, 2016) (arguing why AWS should be banned).

70. *See* Anderson et al., *supra* note 19, at 392–93 (discussing whether to apply the law of war to uses of force or to systems).

to hold states responsible for the actions of their AWS[71] or create a regulatory regime might depend in part on the task the weapon is doing.

Determining how to think about the introduction of new military technologies, and ways of using those technologies, is always complicated.[72] From the crossbow to the submarine to nuclear weapons, states and international organizations have grappled with questions concerning the legality and ethics of new weapon systems. The reality of any potential autonomous weapon systems will almost certainly be more mundane than images from the movies of robots seeking to destroy humanity. That being said, AWS could still raise important questions that will be necessary for scholars and policymakers to grapple with to ensure that the use of military force continues to occur in concert with the law of war. Focusing on the actual activities of potential autonomous weapon systems in conflict situations is one way to move the discussion forward and help everyone better assess how to ensure that accountable and responsible humans remain at the center of all decisions concerning war and peace.

71. *See* Rebecca Crootof, *War Torts*, 164 U. PA. L. REV. (forthcoming 2016) (highlighting the need to develop "war torts" to assure liability regarding crimes committed by AWS).

72. *See* MICHAEL C. HOROWITZ, THE DIFFUSION OF MILITARY POWER: CAUSES AND CONSEQUENCES FOR INTERNATIONAL POLITICS 2 (2010) (discussing the importance of the difference between introducing new technology on the battlefield and full integration of that technology into national strategy).