WORKING PAPER | FEBRUARY 2015

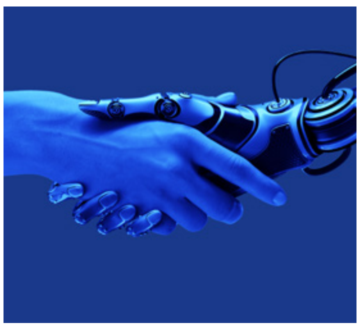

An Introduction to AUTONOMY *in* WEAPON SYSTEMS

By: Paul Scharre and Michael C. Horowitz

Center for a New American Security

About the Authors

Michael C. Horowitz is an Adjunct Senior Fellow at CNAS and an Associate Professor of Political Science at the University of Pennsylvania.

Paul Scharre is a Fellow and Director of the 20YY Warfare Initiative at CNAS.

The Ethical Autonomy project is a joint endeavor of CNAS' Technology and National Security Program and the 20YY Warfare Initiative, and is made possible by the generous support of the John D. and Catherine T. MacArthur Foundation.

PREFACE

Information technology is driving rapid increases in the autonomous capabilities of unmanned systems, from self-driving cars to factory robots, and increasingly autonomous unmanned systems will play a significant role in future conflicts as well. "Drones" have garnered headline attention because of the manner of their use, but drones are in fact remotely piloted by a human, with relatively little automation and with a person in control of any weapons use at all times. As future military systems incorporate greater autonomy, however, the way in which that autonomy is incorporated into weapon systems will raise challenging legal, moral, ethical, policy and strategic stability issues.

Nation-states and activists in the United States and abroad are already debating how advances in autonomy will influence warfare – and what, if anything, should be done. Activists from 54 non-governmental organizations have launched a "Campaign to Stop Killer Robots." In May of 2014, state parties to the United Nations Convention on Certain Conventional Weapons (CCW) held initial discussions on autonomous weapons, and discussions will continue at the April 2015 meeting of the CCW.[1]

Governments and militaries are only beginning to grapple with how to address the challenges and opportunities associated with increased autonomy. Technology is moving fast in this area. Yet few states have guidelines on how autonomy should be included in future weapon systems, with the United States being a notable exception.[2]

The Center for a New American Security (CNAS) has launched a project on Ethical Autonomy, which will examine the legal, moral, ethical, policy and strategic stability dimensions of increased autonomy in future weapon systems. The goal of CNAS' Ethical Autonomy project is to help states, activists, academics and militaries grapple with the challenging issues raised by autonomy in future weapons. This dialogue is necessary to ensure an appropriate balance between ethical and strategic stability considerations, technological opportunities and future warfighting needs.

CNAS' Ethical Autonomy project will result in a series of working papers examining various issues associated with autonomy in weapons, which will ultimately culminate in a final report in late 2015. This first working paper will revolve around framing the issue of autonomy in weapons and providing historical context on the use of autonomy to date. This paper aims to both introduce the topic and clarify the existing debate to help achieve a common understanding of the technology. Subsequent papers will delve deeper into specific legal, moral, ethical, policy or strategic stability issues raised by autonomous weapon systems.

1. The United Nations Office at Geneva, "The Convention on Certain Conventional Weapons," www.unog.ch/80256EE600585943/(httpPages)/4F0DEF093B4860B4C1257180004B1B30?OpenDocument; The United Nations Office at Geneva, "The Convention on Certain Conventional Weapons: Latest Information," http://www.unog.ch/80256EE600585943/%28httpPages%29/3CFCEEEF52D553D5C1257B0300473B77?OpenDocument.
2. United States Department of Defense, Autonomy in Weapon Systems, 3000.09 (November 21, 2012), www.dtic.mil/whs/directives/corres/pdf/300009p.pdf..

I. INTRODUCTION

The past several years have seen an explosion in commentary and scholarship on the topic of autonomous weapons.[3] The prospect that increased autonomy might change or diminish the control of humans over the use of force has garnered the attention of activists, lawyers, political scientists, ethicists, philosophers, and military and national security professionals, rapidly generating a growing body of work.[4]

It is crucial to understand how today's weapons already incorporate autonomy. Arguments about whether to incorporate autonomy into future weapons are sometimes made from the perspective of an idealized version of human control, divorced from the realities of war and how existing weapons use forms of autonomy.

In fact, autonomy is already used for a wide variety of military tasks, including many related to the use of force. These include: identifying, tracking, prioritizing and cueing targets; deciding the timing of when to fire a weapon; maneuvering and homing in on targets; and detonation timing. According to research conducted for this paper, at least 30 countries have defensive systems with human-supervised autonomous modes that are used to defend military bases and vehicles from short-warning attacks, where the time of engagement would be too short for a human to respond.[5] And, in a few rare exceptions, autonomy is used for weapons to select and engage targets on their own.

Moreover, there is no internationally agreed-upon definition of what constitutes an "autonomous weapon," making clear communication on the topic more difficult. The United States Department of Defense, the International Committee of the Red Cross, and the UN Special Rapporteur for extrajudicial, summary or arbitrary executions all use similar definitions, but there is no standard definition that is universally embraced.[6] Activists campaigning to ban "killer robots" (their term for autonomous weapon systems) have yet to put forward a clear definition of their own or even clarify what, precisely, they are advocating should be banned.[7]

This lack of clarity in terminology is further compounded by the fact that some are calling for autonomous weapons to be regulated or banned even before consensus exists on how to define the category.[8] Thus, at present, definitions are tied up in debates over the technology itself. While some degree of debate over the scope of a weapon's category is intrinsic to negotiations on a potential agreement concerning that weapon, lack of clarity on basic terminology itself is a recipe for disaster.

3. It is most appropriate to think of "weapon systems" that comprise not only the munition itself, but also the launching platform, sensor, targeting process and any other elements such as communications links that are necessary for an engagement to occur, some of which may be distributed across multiple physical platforms. Thus, the more technically accurate term is "autonomous weapon system." For much of this paper, however, for the sake of brevity we will use the shorthand "autonomous weapon." In this paper, "autonomous weapon" and "autonomous weapon system" should be treated as interchangeable, and it should be understood that we are referring to not only the munition but the entirely of the weapon system.

4. For a comprehensive bibliography on the topic of autonomous weapons, see "The Ethical Autonomy Project: Bibliography" (Center for a New American Security), www.cnas.org/research/defense-strategies-and-assessments/20YY-Warfare-Initiative/Ethical-Autonomy/bibliography.

5. See Appendix B.

6. See Appendix A.

7. The official page for the Campaign to Stop Killer Robots is available at: www.stopkillerrobots.org/.

8. Colin Allen and Wendell Wallach, "Framing Robot Arms Control," Ethics and Information Technology (2013), link.springer.com/article/10.1007%2Fs10676-012-9303-0; Peter Asaro, "On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making," International Review of the Red Cross, 94 no. 886 (Summer 2012), 687-709, www.icrc.org/eng/assets/files/review/2012/irrc-886-asaro.pdf; "Autonomous Weapons," Article36.org, www.article36.org/issue/weapons/autonomous-weapons/; "Losing Humanity: The Case Against Killer Robots" (Human Rights Watch), 2012, www.hrw.org/sites/default/files/reports/arms1112ForUpload_0_0.pdf.

This is particularly difficult since there are few, if any, obvious examples of "autonomous weapons." Activists have cited some weapons today as "precursors" to potential autonomous weapons in the future, but concerns about autonomous weapons are usually aimed at potential future weapons and the general trend of increasing automation, not at any specific weapons that already exist.[9] Thus, one starting point to gain purchase on the definition of autonomous weapons is the understanding that, whatever they are defined to be, they generally do not exist today. Therefore, while there may be rare exceptions, defining "autonomous weapons" in a way that captures wide swathes of existing weapons is almost certainly wrong, since it misses the essence of what future, increased autonomy might bring.

In the interests of helping to clarify the debate, this paper addresses three key questions:

- What is autonomy?
- How is autonomy used in weapons today?
- What is an "autonomous weapon"?

This paper does not examine in detail the legal, moral, ethical and policy issues raised by increased autonomy in future weapons. Rather, the purpose of this paper is to help clarify, as a prerequisite to examining legal, moral, ethical and policy issues, what an autonomous weapon is, how autonomy is already used, and what might be different about increased autonomy in the future.

The intent is that this paper will not only help lay the groundwork for follow-on research as part of CNAS' Ethical Autonomy project, but also serve as a launching point for others writing on the topic of autonomous weapons. In the span of only a few years, there has been an explosion of interest and concern on the topic of autonomous weapons, and at this juncture taking a step back to review and frame the issue can help focus future debate. The rapid march of technological progress in autonomy, driven largely by civilian advances in robotics and computer processing, underscores a sense of urgency in understanding more clearly the challenges posed by potential future autonomous weapons.

II. WHAT IS "AUTONOMY?"

Different people use the word "autonomy" in different ways, making communicating about trends involving autonomy in weapons particularly challenging. Even setting aside the notion of weapons for a moment, the term "autonomous robot" conjures up wildly different images, ranging from a household Roomba to the sci-fi Terminator.[10] Writers or presenters on the topic of autonomy often articulate "levels of autonomy," but their levels rarely match up, leading a recent U.S. Defense Science Board report on autonomy to throw out the concept of "levels" of autonomy altogether.[11]

9. "Ban 'Killer Robots' Before It's Too Late" (Human Rights Watch), November 19, 2012, www.hrw.org/news/2012/11/19/ban-killer-robots-it-s-too-late. John Markoff, "Fearing Bombs That Can Pick Whom to Kill," The New York Times, November 11, 2014, www.nytimes.com/2014/11/12/science/weapons-directed-by-robots-not-humans-raise-ethical-questions.html?_r=0; "Missile Systems and Human Control" (Campaign to Stop Killer Robots), November 24, 2014, www.stopkillerrobots.org/2014/11/missile-systems-and-human-control/.
10. This section is excerpted from Paul Scharre, "What is Autonomy?," War on the Rocks, (forthcoming).
11. Defense Science Board, Task Force Report: The Role of Autonomy in DoD Systems (July 2012), www.acq.osd.mil/dsb/reports/AutonomyReport.pdf.

However one defines autonomy, interest in it is growing among militaries around the world.[12] In the fall of 2014, U.S. Undersecretary of Defense for Acquisition, Technology and Logistics, Frank Kendall, announced the commissioning of a new study focused on “the science, engineering, and policy problems that must be solved to permit greater operational use of autonomy across all warfighting domains.”[13] NATO is attempting to better understand the military utility of autonomous systems.[14] South Korea has developed robotic sentries with “automatic surveillance” to monitor the demilitarized zone with North Korea.[15] And Russia is developing ground sentry robots to guard missile sites, although their degree of autonomy is unclear.[16]

In its simplest form, autonomy is the ability of a machine to perform a task without human input. Thus, an “autonomous system” is a machine, whether hardware or software, that, once activated, performs some task or function on its own.

Autonomous systems are not limited to robots or uninhabited vehicles. In fact, autonomous, or automated, functions are included on equipment that people use every day. Most cars today include anti-lock brakes, traction and stability control, power steering, emergency seat belt retractors and air bags. Higher-end cars may include intelligent cruise control, automatic lane keeping, collision avoidance and automatic parking. For military aircraft, automatic ground collision avoidance systems (auto-GCAS) can take control of a human-piloted aircraft if a pilot becomes disoriented and is about to fly into terrain. And modern commercial airliners have a high degree of automation available throughout every phase of a flight. Increased automation or autonomy can have many advantages, including increased safety and reliability, improved reaction time and performance, reduced personnel burden with associated cost savings, and the ability to continue operations in communications-degraded or -denied environments.

Parsing out how much autonomy a system has is important for understanding the challenges and opportunities associated with increasing autonomy. There is a wide gap, of course, between a Roomba and a Terminator. Rather than search in vain for a unified framework of “levels of autonomy,” a more fruitful direction is to think of autonomy as having three main axes, or dimensions, along which a system can vary. These dimensions are independent, and so autonomy does not exist on merely one spectrum, but on three spectrums simultaneously.

The Three Dimensions of Autonomy

What makes understanding autonomy so difficult is that autonomy can refer to at least three completely different concepts:

- The human-machine command-and-control relationship

12. For a discussion of some of the reasoning behind this growing interest in autonomy, see Kenneth Anderson, Daniel Reisner, and Matthew C. Waxman, “Adapting the Law of Armed Conflict to Autonomous Weapon Systems,” International Law Studies, 90 (September 2014), papers.ssrn.com/sol3/papers cfm?abstract_id=2477095 and Kenneth Anderson and Matthew C. Waxman, “Law and Ethics for Autonomous Weapons: Why a Ban Won’t Work and How the Laws of War Can,” Columbia University Public Law & Legal Theory Research Paper Series, Research Paper No. 13-351, April 2013, papers.ssrn.com/sol3/papers.cfm?abstract_id=2250126.

13. Frank Kendall, Undersecretary of Defense for Acquisition, Technology and Logistics, “Terms of Reference – Defense Science 2015 Summer Study on Autonomy,” Memorandum for Chairman, November 17, 2014, 1, www.acq.osd.mil/dsb/tors/TOR-2014-11-17-Summer_Study_2015_on_Autonomy.pdf.

14. Multinational Capability Development Campaign, Role of Autonomous Systems in Gaining Operational Access (October 29, 2014), innovationhub-act.org/sites/default/files/u4/Policy%2520Guidance%2520Autonomy%2520in%2520Defence%2520Systems%2520MCDC%25202013-2014%2520final.pdf.

15. Mark Prigg, “Who goes there? Samsung unveils robot sentry that can kill from two miles away,” Daily Mail, September 15, 2014, www.dailymail.co.uk/sciencetech/article-2756847/Who-goes-Samsung-reveals-robot-sentry-set-eye-North-Korea.html.

16. David Hambling, “Armed Russian robocops to defend missile bases,” New Scientist, April 23, 2014, www.newscientist.com/article/mg22229664.400-armed-russian-robocops-to-defend-missile-bases.html.

- The complexity of the machine
- The type of decision being automated

These are all important features of autonomous systems, but they are different ideas, and often people tend to mix them together.

THE HUMAN-MACHINE COMMAND-AND-CONTROL RELATIONSHIP

The first way in which the word autonomy is used refers to the relationship between the person and the machine. Machines that perform a function for some period of time, then stop and wait for human input before continuing, are often referred to as “semiautonomous” or as having a “human in the loop.” Machines that can perform a function entirely on their own but have a human in a monitoring role, with the ability to intervene if the machine fails or malfunctions, are often referred to as “human-supervised autonomous” or “human on the loop.” Machines that can perform a function entirely on their own with humans unable to intervene are often referred to as “fully autonomous” or “human out of the loop.” In this sense, “autonomy” is not about the intelligence of the machine, but rather its relationship to a human controller.

THE COMPLEXITY OF THE MACHINE

The word “autonomy” can also be used in a completely different way to refer to the complexity of the system. Regardless of the human-machine command-and-control relationship, words such as “automatic,” “automated” and “autonomous” are often used to refer to a spectrum of complexity of machines. The term “automatic” is often used to refer to systems that have very simple, mechanical responses to environmental input. Examples of “automatic” devices in the civilian world include toasters and mechanical thermostats. In the military world, trip wires and mines can be described as “automatic” devices. The term “automated” is often used to refer to more complex, rule-based systems. Self-driving cars and modern programmable thermostats are examples of such systems. Sometimes the word “autonomous” is reserved for machines that execute some kind of self-direction, self-learning or emergent behavior that is not directly predictable from an inspection of its code. An example would be a self-learning robot that taught itself how to walk or the Nest “learning thermostat.”[17]

Others will reserve the word autonomous only for entities that have intelligence and free will, but these concepts hardly add clarity. “Artificial intelligence” is a loaded term that can refer to a wide range of systems, from those that exhibit near-human or super-human intelligence in a narrow domain, such as playing chess (Deep Blue), playing Jeopardy (Watson) or programming subway repair schedules, to potential future systems that might have human or super-human general intelligence.[18] But whether general intelligence leads to free will, or whether humans even have free will, is itself debated.

What is particularly challenging is that there are no clear boundaries between these degrees of complexity, from “automatic” to “automated” to “autonomous” to “intelligent,” and different people may disagree on what to call any given system.

---

17. Hod Lipson, “Building ‘self-aware’ robots,” TED, March 2007, www.ted.com/talks/hod_lipson_builds_self_aware_robots; Nest, “Life with Nest Thermostat,” https://nest.com/thermostat/life-with-nest-thermostat/.

18. On Deep blue, see IBM, “Icons of Progress: Deep Blue,” www-03.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/. On Watson, see IBM, “Icons of Progress: A Computer Called Watson,” www-03.ibm.com/ibm/history/ibm100/us/en/icons/watson/. Also see Hal Hodson, “The AI boss that deploys Hong Kong’s subway engineers,” New Scientist, July 4, 2014, www.newscientist.com/article/mg22329764.000-the-ai-boss-that-deploys-hong-kongs-subway-engineers.html#.VLqYckfF_X4; James Barrat, Our Final Invention: Artificial Intelligence and the End of the Human Era (New York: Thomas Dunne Books, 2013).

TYPE OF FUNCTION BEING AUTOMATED

Ultimately, it is meaningless to refer to a machine as "autonomous" or "semi-autonomous" without specifying the task or function being automated. Different decisions have different levels of complexity and risk. A mine and a toaster have radically different levels of risk, even though both have humans "out of the loop" once activated and both use very simple mechanical switches. The task being automated, however, is much different. Any given machine might have humans in complete control of some tasks and might autonomously perform others. For example, an "autonomous car" may drive from point A to point B on its own, but a person still chooses the final destination, and potentially even the route. In that case, the "autonomous car" is only autonomous with respect to some functions.

Even in tasks relating to the use of military force, there are different engagement-related functions, not all of which are equally significant when it comes to thinking about the role of human control. Engagement-related tasks include, but are not limited to, acquiring, tracking, identifying and cueing potential targets, aiming weapons, selecting specific targets for engagement, prioritizing targets to be engaged, timing of when to fire, maneuvering and homing in on targets, and the detonation itself.

Parsing Autonomy

There is not a single spectrum along which autonomy moves. The paradigm of human versus machine is a common science fiction meme, but a better framework would be to ask which tasks a person does, and which by a machine. A recent NATO policy document came to a similar conclusion, recommending a framework of thinking about "autonomous functions" of systems, rather than characterizing an entire vehicle or system as "autonomous."[19]

Importantly, these three dimensions of autonomy are independent. Thus, the intelligence, or complexity, of the machine is a separate concept from its ability to perform a task on its own. Increased intelligence or more sophisticated machine reasoning to perform a task does not necessarily equate to transferring control over a wider range of tasks from the human to the machine. Similarly, the human-machine command-and-control relationship differs from the complexity of the machine or the tasks being performed. A thermostat functions on its own without any human supervision or intervention when you leave your house, but it still has a limited set of functions it can perform.

Therefore, when talking about "autonomous weapons," we must ask: "Autonomous" with respect to which functions or tasks? Which tasks or functions are the critical ones that, when performed by a machine instead of a human, make it an "autonomous weapon?"

First, however, we must start with an understanding of how autonomy is already used in weapons today. If some existing weapon systems are precursors, but autonomous weapons generally do not exist today, then a better understanding of how exactly autonomy is already used – and how it has generally not been used to-date – can help us gain purchase on the question of how to define an autonomous weapon.

19. Multinational Capability Development Campaign, Role of Autonomous Systems in Gaining Operational Access.

III. HOW IS AUTONOMY USED IN WEAPONS?

Various forms of autonomy have been used in military systems for over seventy years. The first homing munitions date back to World War II, and automated defensive systems with human-supervised modes have been in existence for decades and are used by many militaries around the globe.[20] Automation is not merely important, but essential for modern militaries to conduct many tasks, including identifying targets by radar or delivering precision-guided weapons.

Since nearly all of the literature on autonomous weapons focuses on the task of selecting and engaging specific targets, the next section delineates how autonomy is used in weapons today along a similar typology.[21] Specifically, we will divide weapons into three categories:

1. Human "in the loop" for selecting and engaging specific targets – Weapon systems that use autonomy to engage individual targets or specific groups of targets that a human has decided are to be engaged.
2. Human "on the loop" for selecting and engaging specific targets – Weapon systems that use autonomy to select and engage targets where a human has not decided those specific targets are to be engaged, but human controllers can monitor the weapon system's performance and intervene to halt its operation if necessary.
3. Human "out of the loop" for selecting and engaging specific targets – Weapon systems that use autonomy to select and engage targets where a human has not decided those specific targets are to be engaged, and human controllers cannot monitor the weapon system's performance and intervene to halt its operation if necessary.

Human "in the loop" for selecting and engaging specific targets

World War II saw the advent of homing munitions that used automation to guide projectiles onto their targets. One of the first, the German *Zaunkönig* (Wren) torpedo, had a passive acoustic homing seeker. Once launched, it would run straight for 400 meters, and then activate its homing seeker, which consisted of two hydrophone receivers. These receivers would sense the noise of a ship's propeller in the water and then steer the torpedo toward the ship, ultimately destroying it. The Wren was the first of an entirely new class of weapon: the guided munition.[22]

GUIDED MUNITIONS

Guided munitions are "projectiles, bombs, missiles, torpedoes and other weapons that can actively correct for initial-aiming or subsequent errors by homing on their targets or aim-points after being fired, released, or launched."[23] There are many types of guided munitions, but what they all have in common is that autonomy assists the weapon in striking an individual target or a specific group of targets that a person has decided should be engaged. The person launching the weapon knows what specific targets are to be engaged, and has made a conscious decision that those targets should be destroyed.

---

20. See Appendix B. Also see Barry D. Watts, "Six Decades of Guided Munitions and Battle Networks: Progress and Prospects" (Center for Strategic and Budgetary Assessments, 2007) and David A. Koplow, Death by Moderation: The U.S. Military's Quest for Useable Weapons (New York: Cambridge University Press, 2012).
21. See "The Ethical Autonomy Project: Bibliography" (Center for a New American Security).
22. "TV (G7es) Acoustic Homing Torpedo," German U-Boat, www.uboataces.com/torpedo-tv.shtml; , Naval Weapons of Germany (New York: LLC Books, 2010).
23. Watts, "Six Decades of Guided Munitions and Battle Networks: Progress and Prospects."

The Wren acoustic homing torpedo was the first guided munition, but all of the major powers in World War II conducted research on guided munitions, and over subsequent decades they expanded into other domains, including air and ground warfare, and matured considerably. Today, modern militaries around the globe employ guided weapons in the form of air-to-air homing missiles, anti-ship missiles, cruise missiles or air-to-ground guided missiles. The guidance mechanisms of these projectiles vary significantly, but broadly speaking they can be broken into two categories: projectiles designed to hit a particular target based on its signature (go-onto-target) and projectiles designed to hit a particular geographic location where the target is located (go-onto-location-in-space).

Weapons designed to hit a particular target (go-onto-target) are conceptually similar to the Wren torpedo; they home in on moving targets with a particular signature. Some use passive sensors, merely sensing signals from the environment such as noise, heat or electromagnetic signals and adjusting direction as needed to close in on a target. Others have active sensors that send out signals, such as radar, and look for a target's return. Some "lock on" to the target before launch. Others are launched at a target, like the Wren, then travel a distance and activate their seeker in order to more accurately home in on the target in the endgame, "locking on" after launch. This would be the case if the onboard seeker on the munition did not have sufficient range to detect the target at the distance it was launched. Over time, seeker processors have become increasingly sophisticated as militaries have competed in a constant race to field countermeasures to confuse and distract adversary seekers, while militaries have simultaneously improved the ability of seekers to interpret signals in the environment and distinguish countermeasures from true targets.[24]

In contrast, go-onto-location-in-space guided munitions are not designed to hit a particular target based on its signature, but rather a particular point in geographic space where a target is known to be located. For example, GPS-guided bombs fly to a particular point based on signals from the satellite-based global positioning system. Cruise missiles frequently use terrestrial guidance systems that sense the ground beneath the missile, typically via a radar altimeter, and compare it to an onboard map, allowing the missile to fly a pre-programmed route. Laser-guided bombs fly toward a spot designated by a laser pointer, either from an aircraft or troops on the ground. In addition, submarine-launched ballistic missiles use celestial guidance, or star-matching, for navigation.[25]

Guided munitions also vary in other important ways:

- Some guided munitions have the ability to be controlled, aborted or retargeted in flight. For example, the latest version of the Tomahawk land-attack cruise missile, the Block IV, has a communications link that allows in-flight retargeting.[26]
- Some guided munitions cannot be retargeted in flight and are considered "fire and forget" weapons; once launched, they cannot be recalled.[27] Because of the time of flight, the time between weapon launch and

24. One way this is done is by employing multiple sensing methods on a seeker, making it a "multi-mode" seeker. Whitney, Bradley & Brown, Inc., "Multi-Mode Precision Strike Weapons," October 17, 2006, www.dtic.mil/ndia/2006psa_psts/kuz.pdf. For a recent example, see discussions of Raytheon's multi-mode seeker upgrades to the Tomahawk missile. Christopher P. Cavas, "Raytheon Working on Tomahawk With Seeker," Defense News, January 14, 2015, www.defensenews.com/story/defense/show-daily/surface-navy-assn/2015/01/13/raytheon-working-on-tomahawk-with-active-seeker/21718359/.
25. Craig Payne, Principles of Naval Weapons Systems (Annapolis, MD: Naval Institute Press, 2006), p. 309.
26. NAVAIR Public Affairs, "Tactical Tomahawk Takes Flight," United States Navy, August 27, 2002, www.navy.mil/submit/display.asp?story_id=3311.
27. "Fire and forget" has, of course, been the norm for projectiles in war since the sling and stone. It is only in very recent modern times with the invention of precision-guided weapons with communications links that humans have had the ability to retarget or abort a projectile while in flight.

striking a target could be significant for some types of weapons, up to several hours for cruise missiles.[28] This means that human controllers may not have real-time, high-fidelity awareness of the target immediately before the moment of impact. Instead, decisions about the use of force are based on the nature of the target and a general understanding of the likelihood of potential collateral damage or civilian casualties, given the target's location.[29]

- In some cases, multiple guided weapons might be launched at a group of targets. For example, an aircraft pilot might launch multiple air-to-air guided missiles at a group of oncoming enemy aircraft.
- In other cases, some weapons have multiple guided sub-munitions. For example, a sensor fuzed weapon releases smaller sub-munitions, each of which have onboard seekers to home in on a specific group of enemy tanks.[30]

As autonomy becomes more sophisticated, it may be possible for a weapon to identify a specific tank, ship or airplane because of a unique signature of that particular vehicle. In general, however, guided munitions are used to strike only the specific target or targets that a person has chosen because the freedom of maneuver of the munition is fairly constrained. Go-onto-location-in-space weapons go to a particular place chosen by a human. Go-onto-target weapons that "lock on" before launch are already locked onto a specific target chosen by the human controller before they are released. Go-onto-target weapons that "lock on" after launch are constrained in time and space in order to ensure they will engage only the target chosen by the human controller. They are not used to search over a wide area for a target or loiter until a target arrives, but rather are launched in a particular direction at a specific target. The field of view and loiter time of the seeker is relatively constrained, with the intent of ensuring that only the desired target is within the seeker's field of view when the seeker activates.[31]

Consider the example of the Wren acoustic homing torpedo. World War II submariners did not simply launch the Wren wildly into the deep blue of the ocean to go searching for ships. To do so would have been militarily ineffective, wasting a torpedo. A submarine would patrol and, once it found an enemy ship, launch the torpedo at the ship. At this stage, the Wren was no different from the types of unguided projectiles, such as rocks and arrows, that have characterized warfare for millennia. After 400 meters, the Wren would activate its acoustic homing seeker, however, which would allow the Wren to more accurately zero in on the enemy ship, striking the desired target.[32] In this sense, the Wren functioned very similar to an unguided projectile in terms of how it was used; it simply had the addition of a sensor that allowed it to correct for aiming errors and home in on targets if it was off-course.

IDENTIFYING, ACQUIRING, TRACKING, CUEING AND PRIORITIZING TARGETS

Modern militaries also use autonomy in several other engagement-related functions prior to weapons release. Radars and other sensors help acquire, track, identify and cue potential targets to human operators.

---

28. See Watts, "Six Decades of Guided Munitions and Battle Networks: Progress and Prospects." This is more likely to be the case for go-onto-location-in-space munitions since, for go-onto-target munitions, the target will likely move over a long enough time period.
29. This is not a new development. Humans have been using weapons where they do not have real-time sight of the target area since at least the invention of the catapult.
30. Textron Defense Systems, "CBU-105 Sensor Fuzed Weapon/BLU-108 Submunition," 2014, www.textronsystems.com/sites/default/files/pdfs/product-info/sfw_brochure-small.pdf; Textron Systems, "Textron Systems - Sensor Fuzed Weapon (SFW) CBU-105 With BLU-108 Submunitions," https://www.youtube.com/watch?v=9HkauulyDsM.
31. This is why it is particularly important to examine not only the technical specifications of the weapon, but also the doctrine and procedures for its use.
32. LLC Books, eds., Naval Weapons of Germany.

For some sophisticated weapons that require precise engagement, the timing of precisely when to fire the weapon and/or when to detonate are automated.[33]

When used in conjunction with guided munitions, current uses of force can therefore involve a high degree of automation in a variety of engagement-related functions, though the decision of which specific target or group of targets are to be engaged is still made by a person.

Consider, for example, an air-to-air missile engagement. In the case of beyond visual range engagements, potential targets are identified by a computer which conveys this information to the pilot; the pilot does not have the ability to visually confirm the identity of the target, nor does the pilot even have a picture image of the target. Instead, the pilot is relying on information fed from a computer, typically based on radar. Based on this data as well as other information such as altitude, airspeed, identification friend or foe (IFF) signals and an understanding of the overall situation, the pilot makes a determination whether it is appropriate to engage this particular aircraft or group of aircraft. If the pilot decides that it is appropriate to engage, then the pilot launches an air-to-air homing missile at the enemy aircraft. If there are multiple aircraft to be engaged in the same area, the pilot might launch multiple missiles at the group of aircraft. If the missiles function with a lock-on-before-launch feature, then specific target data will be passed to the missile before launch. If the missile is functioning in a lock-on-after-launch capacity, then the missile does not have any specific targeting data on the particular enemy aircraft to be engaged. Instead, the missile will fly to a point in space and then activate its seeker, looking for targets that meet its parameters. The pilot ensures that the missile only engages the targets that he or she intends to engage by the use of tactics, techniques and procedures to ensure that, when the missile activates its seeker, the only targets within the seeker's field of view are those that the pilot intends to engage.

Note that in this situation, the missile is "autonomous" in the sense that, once it is fired, the pilot does not have the ability to recall or control the missile. There is also a high degree of automation in numerous stages of the decision cycle. However, a person decides which specific targets to engage. The weapon is not designed to search over a wide area and select its own targets. Instead, the weapon is designed and used to home in on the individual target or specific group of targets that the pilot intends to be engaged.

PRECISION-GUIDED WEAPONS INCREASE HUMAN CONTROL OVER WARFARE

Since the advent of the Wren homing torpedo in World War II, guided munitions have become an increasingly important component of modern militaries. One of the earliest prominent uses of air-to-ground guided munitions by the United States military, for example, was in Operation Linebacker in Vietnam.[34] Their use, alongside automated target identification radars and other autonomous functions, has given humans greater, not lesser, control over the use of force in warfare.[35] This is because they have made weapons more precise, allowing militaries to target particular buildings or even floors of buildings, rather than carpet-bombing whole cities as combatants did before the advent of precision-guided munitions. In fact, some human rights groups have suggested that not using precision-guided weapons when dropping bombs in an urban area could be a war crime.[36]

33. Watts, "Six Decades of Guided Munitions and Battle Networks: Progress and Prospects."
34. See Kenneth P. Werrell, Chasing the silver bullet: U.S. Air Force weapons development from Vietnam to Desert Storm (Washington, DC: Smithsonian Books, 2003).
35. Koplow, Death by Moderation: The U.S. Military's Quest for Useable Weapons.
36. "Ukraine: Unguided Rockets Killing Civilians" (Human Rights Watch), July 24, 2014, www.hrw.org/news/2014/07/24/ukraine-unguided-rockets-killing-civilians.

Next-generation precision-guided weapons continue this trend, leveraging automation to help humans engage specific targets that humans have decided should be engaged. While some activists have raised concerns about these weapons, more sophisticated autonomy in and of itself is not a bad thing.[37] More advanced automated target recognition could be used to help ensure that a munition hits the target selected by the human operator, for example. Rather, the key place to focus attention is which tasks are being automated and which does the human retain. Weapon systems that incorporate autonomy to more accurately or effectively engage specific targets or clusters of targets that have been chosen by a person do not raise substantively new issues, but rather are a continuation of the way in which autonomy has been used for over seventy years.

Human "on the loop" for selecting and engaging specific targets

At least 30 nations use human-supervised defensive systems with greater autonomy, where humans are "on the loop" for selecting and engaging specific targets.[38] To date, these have been used for defensive situations where the reaction time required for engagement is so short that it would be physically impossible for humans to remain "in the loop" and take positive action before each engagement and still defend effectively. Instead, autonomy is used to complete the engagement against incoming threats that meet certain criteria according to pre-programmed rules determined by humans. Human controllers are cognizant of the specific targets being engaged and can intervene to halt the weapon system, but do not make a positive decision to engage specific targets.

Currently deployed defensive, human-supervised autonomous weapon systems include:

- Air and missile defense systems, such as the U.S. ship-based Aegis and the land-based Patriot systems.
- Shorter-range counter-rocket, artillery and mortar systems to defend land bases, such as the Dutch Goalkeeper system.
- Active protection systems for ground vehicles to automatically engage incoming rockets or missiles, such as the French Shark or Russian Arena systems.

At least 30 nations employ or have in development at least one system of this type, including: Australia, Bahrain, Belgium, Canada, Chile, China, Egypt, France, Germany, Greece, India, Israel, Japan, Kuwait, the Netherlands, New Zealand, Norway, Pakistan, Poland, Portugal, Qatar, Russia, Saudi Arabia, South Africa, South Korea, Spain, Taiwan, the United Arab Emirates, the United Kingdom, and the United States.[39]

In all of these cases, automation is used to defend human-occupied bases or vehicles from being overwhelmed by a rapid barrage of missiles or rockets. These systems are an essential component of defending and surviving in modern warfare, and automated missile defense systems can also serve an invaluable role in defending civilian populations from attack.

These systems have, however, been used very narrowly to date. They are employed to defend human-occupied vehicles and installations; they only target objects (missiles, rockets or aircraft), not people; and human

37. See Mark Gubrud, "Killer Robots and Laser-Guided Bombs: A reply to Horowitz & Scharre," 1.0 Human, December 4, 2014, gubrud.net/?p=398; John Markoff, "Fearing Bombs That Can Pick Whom to Kill." For a more positive take on the possibilities offered by autonomy in weapon systems, see Ronald C. Arkin, Governing Lethal Behavior in Autonomous Robots (New York: CRC Press, 2009).
38. See Appendix B
39. See Appendix B.

controllers supervise the system's operation in real-time and can intervene if necessary.[40] Given the short time required for engagements, human controllers may not be able to intervene before any inappropriate engagements occur, but they can terminate further functioning of the system. Moreover, because human controllers are physically co-located with the system, either on a vehicle or a land base, they have physical access and can exercise hardware-level overrides in the event of a software malfunction or cyber attack.

Human "out of the loop" for selecting and engaging specific targets

There are a limited number of existing weapons that have a human fully out of the loop for selecting specific targets to be engaged, such that the weapon itself is selecting the targets. These weapon systems use autonomy to engage general classes of targets in a broad geographic area according to pre-programmed rules, and humans controllers are not aware of the specific targets being engaged.

LOITERING MUNITIONS

One type of weapon that selects and engages targets on its own, without a human "in" or "on the loop" is a loitering munition. Loitering munitions differ from guided munitions in that they are not launched at a specific target. Rather, they are launched into a general area where they will loiter, flying a search pattern looking for targets within a general class, such as enemy radars, ships or tanks. Then, upon finding a target that meets its parameters, the weapon will fly into the target and destroy it. Thus, unlike a homing munition, the human operator launching the weapon does not know the specific target that is to be engaged, only that the weapon will engage targets of a particular class within a broad geographic area.[41]

Examples of experimental loitering munitions include the U.S. low-cost autonomous attack system (LOCAAS), which was designed to target tanks, and Tacit Rainbow, a loitering anti-radar munition.[42] Neither was ever deployed by the United States military. The United States did have in its inventory a loitering anti-ship missile in the 1980s called the Tomahawk Anti-Ship Missile (TASM).[43] It was removed from the U.S. inventory in the 1990s.

The only currently operational loitering munition, according to extensive research, appears to be the Israeli Harpy.[44] The Harpy is an anti-radar weapon that flies a search pattern over a wide area searching for enemy

40. Some sources report a South Korean sentry robot, which targets personnel, as having the capacity to autonomously engage targets in a human-supervised mode. Other sources state that authorization to fire requires a person "in the loop." See "Samsung Techwin SGR-A1 Sentry Guard Robot," globalsecurity.org, updated November 7, 2011, www.globalsecurity.org/military/world/rok/sgr-a1.htm; LCDR Jesse Hilliker, U.S. Navy, "Should We Turn the Robots Loose?," U.S. Naval War College, May 2010, 6-7, www.dtic.mil/dtic/tr/fulltext/u2/a525286.pdf; Tim Hornyak, "Korean machine-gun robots start DMZ duty," CNET, July 14, 2010, www.cnet.com/news/korean-machine-gun-robots-start-dmz-duty/; Henry Ridgwell, "Demand Grows for Ban on 'Killer Robots'," Voice of America, May 24, 2013, www.voanews.com/content/demand-grows-for-global-ban-on-killer-robots/1667710.html.

41. Andrea Gilli and Mauro Gilli, Sr., "The Diffusion of Drone Warfare? Industrial, Infrastructural and Organizational Constraints," April 16, 2014, papers.ssrn.com/sol3/papers.cfm?abstract_id=2425750.

42. Andreas Parsch, "Lockheed Martin LOCAAS," Directory of U.S. Military Rockets and Missiles, September 17, 2006, http://www.designation-systems.net/dusrm/app4/locaas.html. National Museum of the U.S. Air Force, "Northrop AGM-136A Tacit Rainbow," August, 3, 2011, http://www.nationalmuseum.af.mil/factsheets/factsheet.asp?id=418.

43. Carlo Kopp, "Tomahawk Cruise Missile Variants: BGM/RGM/AGM-109 Tomahawk/TASM/TLAM/GCLM/MRASM: Technical Report APA-TR-2005-0702" (Air Power Australia, April 2012), www.ausairpower.net/Tomahawk-Subtypes.html. Despite the name, the Tomahawk Anti-Ship Missile bears no functional relation to the more commonly known Tomahawk Land Attack Missile (TLAM).

44. Israel Aerospace Industries, "Harpy Loitering Weapon," www.iai.co.il/2013/16143-16153-en/IAI.aspx. Israel has sold the Harpy to Turkey, South Korea, India, and China. See "Harpy Air Defense Suppression System," Defense Update, updated March 4, 2006, defense-update.com/directory/harpy.htm. China is alleged to have reverse-engineered its own variant. See "UAV/UCAV," Chinese Military Aviation, May 7, 2014, chinese-military-aviation.blogspot.ch/p/uav.html.

radars. If it detects any radar that meets its criteria, the Harpy then dive-bombs into the radar, destroying it.[45]

In a technical sense, loitering munitions chiefly differ from guided munitions in terms of their time of flight and the geographic area they cover. The sensor and algorithm used for a loitering munition may not be any more sophisticated than that of a guided munition. A lock-on-after-launch guided munition may similarly only be able to recognize a particular class of targets. The differences in terms of loiter time and geographic area result in a significant difference in how they are used, however, and in the role of the human controller in terms of selecting specific targets.

Because the freedom of a guided munition is limited in both time and space, guided munitions are only useful if launched at specific targets. Otherwise, if the guided munition activates its seeker and there is not a target directly in its field of view, the munition will likely run out of fuel and fail since it cannot search a wide area to find a target. Loitering munitions, on the other hand, are designed to search over a wide area to find a target. This is a critical distinction: for guided munitions, the human controller must know the specific target to be engaged – otherwise the munition is useless.[46] Loitering munitions are launched without a particular target, or particular target location, in mind.

Similarly, compare a loitering munition to the sensor fuzed weapons described above. While the sensor fuzed weapon deploys multiple homing sub-munitions against a group of tanks, the weapon as a whole is launched against a specific group of enemy vehicles located in a particular area. In fact, it would be difficult to use sensor fuzed weapon any other way, since the time of flight of the weapon is extremely short. If a sensor fuzed weapon were launched into a general area where there was not knowledge of specific enemy tanks then, just like other homing munitions, it would be useless because sensor fuzed weapon cannot loiter or search over a geographic area of several kilometers. For sensor fuzed weapon, therefore, even though it can attack multiple targets simultaneously, the human launching the weapon must know exactly which specific group of vehicles is to be attacked, or the weapon is useless. This highlights the importance of looking not only at the sophistication of the weapon's seeker, but the entirety of the weapon system and its intended use.

ENCAPSULATED TORPEDO MINES

Mines are generally excluded from discussions about autonomous weapons. While there is clearly no person in the loop when the mine detonates, mines are extremely simple devices, and therefore predictable, and a person often emplaces the mine at a particular place, so the location of the mine is known. With remotely delivered mines, the precise location of each mine is not known, but the mined area is known and the mines are immobile. Thus, human control of mines is different than with other weapons, but the mine itself still has very limited freedom of action.[47] A human may not know the specific target that will detonate the mine,

---

45. The Harpy is frequently referred to as an unmanned aerial vehicle (UAV) or unmanned combat air vehicle (UCAV), but the Harpy is essentially a wide-area loitering cruise missile, so it should not be grouped with recoverable drones like the remotely-piloted unmanned aircraft flown by many nations.

46. For example, a Harpy with a maximum range of 500 kilometers and a top speed of 185 kilometers per hour has a time of flight of at least 2.7 hours. This allows it to loiter over a wide area to search for targets. A high-speed anti-radiation missile (HARM), on the other hand, similarly is used to destroy enemy radars but is a homing munition and must be launched at a specific enemy radar in order to be useful. With a maximum range of 90 kilometers and a top-speed of over 1,200 kilometers per hour, the HARM is only airborne for approximately four and a half minutes. Thus, it would not be practicable to use a HARM as a loitering wide-area weapon. See Robert O'Gorman and Chriss Abbott, "Remote Control War: Unmanned combat air vehicles in China, India, Israel, Iran, Russia, and Turkey" (Open Briefing, September 2013), 75; and United States Navy, "AGM-88 HARM Missile," February 20, 2009, http://www.navy.mil/navydata/fact_display.asp?cid=2200&tid=300&ct=2.

47. Accurate record keeping and maintenance of control over the territory where the mine is located is also necessary to maintain control over the mine.

but the human controller will know the specific location.[48]

Encapsulated torpedo mines are a special case, however.[49] Encapsulated torpedo mines are a type of sea mine that, when activated by a passing ship, instead of exploding, open a capsule that then releases a torpedo that engages a target. In this case, the torpedo is not being used to home in on a target that has been selected by a person. Nor is the mine simply blowing up in place. Instead, the mine is given a greater freedom of maneuver to release a torpedo that will then track onto a nearby target. De facto, the mine is selecting and engaging targets on its own. Thus, encapsulated torpedo mines are, arguably, more similar to loitering munitions than to traditional mines.

Encapsulated torpedo mines were part of the U.S. inventory for a number of years, but have since been retired. Russia and China have an encapsulated torpedo mine, the PMK-2, in service today.[50]

IV. WHAT IS AN "AUTONOMOUS WEAPON?"

As previously mentioned, there is no internationally agreed-upon definition for an "autonomous weapon," which complicates discussion of the topic. From a common sense perspective, it makes sense to define an "autonomous weapon" as one that selects and engages targets on its own, and a "semi-autonomous weapon" as one where a human is selecting the specific targets for engagement. These are the definitions most commonly used in the literature on autonomous weapons and are appealing for a number of reasons.

First, the common sense perspective conforms to the widely held assumption that "autonomous weapons" (whatever they are defined to be) do not generally exist today, with some isolated exceptions.[51] Second, it seems to capture the essence about what is so chilling about the vision of a futuristic hunter-killer robot deciding who to kill on its own, without human authorization. Third, it essentially mirrors the definitions already used by a diverse set of organizations that have weighed in on autonomous weapons, including the U.S. Department of Defense, Human Rights Watch, the UN Special Rapporteur on extrajudicial, summary or arbitrary executions, and the International Committee of the Red Cross (ICRC).[52]

Common sense definitions may not be sufficient, however, due to interest in setting exactly where the line is between autonomous and semi-autonomous weapon systems. This type of line is critical for all stakeholders in discussions of autonomy – for military planners attempting to figure out how to direct research and development programs, as well as for activists thinking about what types of activities they wish to place under greater scrutiny.

Additionally, the most likely near-term candidates for autonomous weapons are not malevolent, sentient

48. From a certain perspective, mines could be thought of as a loitering go-onto-location-in-space munition.
49. "Mine Warfare Trends," Mine Warfare Association, May 10, 2011, www.minwara.org/Meetings/2011_05/Presentations/tuespdf/WMason_0900/MWTrends.pdf.
50. Scott C. Truver, "Taking Mines Seriously: Mine Warfare in China's Near Seas," Naval War College Review, 65 no. 2 (Spring 2012), https://www.usnwc.edu/getattachment/19669a3b-6795-406c-8924-106d7a5adb93/Taking-Mines-Seriously--Mine-Warfare-in-China-s-Ne.
51. The defensive, human-supervised autonomous weapon systems in use by at least 30 nations today are a notable exception, although they are used in a very restricted context, with real-time human supervision, and physical access to disable the system if necessary. Human Rights Watch has referred to these defensive systems as "precursors" of the types of weapons they fear, which this definition seems to accurately capture. The Harpy loitering munition and PMK-2 encapsulated torpedo mine are isolated exceptions of fully autonomous weapon systems under this definition.
52. See Appendix A.

humanoid robots (like those of movie fame), but rather something more like wide-area search-and-destroy loitering munitions. Thus, the definitions must clearly distinguish, in a way that is technically rigorous, between autonomous weapons and the precision-guided homing munitions that have been in use for over seventy years.[53]

From a more technical perspective, the following definitions offer a clearer way forward:

- An **autonomous weapon system** is a weapon system that, once activated, is intended to select and engage targets where a human has not decided those specific targets are to be engaged.
- A **human-supervised autonomous weapon system** is a weapon system with the characteristics of an autonomous weapon system, but with the ability for human operators to monitor the weapon system's performance and intervene to halt its operation, if necessary.
- A **semi-autonomous weapon** is a weapon system that incorporates autonomy into one or more targeting functions and, once activated, is intended to only engage individual targets or specific groups of target that a human has decided are to be engaged.

These definitions are intended to clarify and sharpen the distinction between autonomous and semi-autonomous weapons in two important ways:

First, the definition of an autonomous weapon system is intended to clearly delineate that autonomous weapon systems select and engage targets where a human has not decided those specific targets are to be engaged. This essentially spells out what is meant by saying an autonomous weapon would select and engage targets "on its own." At the risk of redundancy, this language's focus on the role of the human operator is designed to clarify any confusion surrounding guided munitions, particularly those that use seekers to home in on targets or specific groups of targets that a human has decided are to be engaged.

Second, the idea of a human decision is embedded within each of the above definitions. This highlights what is different about an autonomous weapon from the perspective of a human operator. The decision to place an autonomous weapon into operation versus a semi-autonomous one is a very different decision. Even in the case of a fire-and-forget homing missile, while the person launching the missile may not be able to recall it, the human operator has made a decision about which individual target or specific group of targets is to be engaged by that homing missile. In the case of an autonomous weapon, however, the human has decided to launch a weapon to seek out and destroy a general class of targets over a wide area, but is not making a decision about which specific targets are to be engaged. Both definitions, however, focus on the decision the human is making or not making and do not apply the word "decision" to something the weapon itself is doing, which could raise murky issues of machine intelligence and free will.

53. We think the distinction is clear in the commonly used "select and engage" definition, but some have seized on the notion that a homing munition has a seeker to suggest that it is "selecting" targets on its own. If one were to define autonomous weapons this way, then the logical conclusion would be that autonomous weapons have been around for nearly three-quarters of a century, so this entire discussion is a lot of fuss for nothing. We think that there is a significant difference between a homing munition that is used to engage a target that has been chosen by a person and a wide-area search-and-destroy weapon that would select and engage targets on its own, so this definition is modified to further clarify that distinction. For more on this topic, see Mark Gubrud, "Killer Robots and Laser-Guided Bombs: A reply to Horowitz & Scharre"; John Markoff, "Fearing Bombs That Can Pick Whom to Kill"; Rebecca Crootof, "The Killer Robots Are Here: Legal and Policy Implications," Cardozo Law Review, 36 (2015), papers.ssrn.com/sol3/papers.cfm?abstract_id=2534567; Michael C. Horowitz and Paul Scharre, "Do Killer Robots Save Lives?" Politico Magazine, November 19, 2014, www.politico.com/magazine/story/2014/11/killer-robots-save-lives-113010.html#.VHODj1XOKY8; "Missile Systems and Human Control" (Campaign to Stop Killer Robots).

It should be noted that because of the different ways that the word autonomy is used, there are many ways a semi-autonomous weapon system could become "more autonomous" in a general sense and still not be an "autonomous weapon system." A semi-autonomous weapon could incorporate greater autonomy into other functions, such as being able to avoid threats or change its route, without changing human control over the decision for which specific targets are to be engaged. Also, a weapon system could become "smarter," incorporating more capable sensors and seekers, but still retain human control over the selection of specific targets. Thus, a semi-autonomous weapon that had automatic target recognition technology and used that technology to ensure that it only engaged individual targets or specific groups of targets that a person had decided are to be engaged would still be a semi-autonomous weapon. This reinforces the importance of looking not only at the weapon's seeker but the entire capability of the weapon system and its intended use.

An alternative terminology for understanding how autonomy is used in weapons – and which functions are most important – could be to focus on the function being automated. In this case, autonomous weapons would be "self-targeted weapons," indicating that the weapon itself is the one choosing the specific target to be engaged. Semi-autonomous weapons would be "human-targeted weapons," making clear that, while other functions may use autonomy, the human is still choosing the weapon. This alternative terminology may be clearer and could help avoid some of the confusion stemming from the word "autonomy."

V. CONCLUSION

Rapid advances in information technology means that capabilities that may seem at present to be distant or impossible may come to fruition sooner than we realize. Activist groups that have catalyzed this debate are to be commended for raising awareness on an important issue. However, leaping to the conclusion that a legally-binding ban on autonomous weapons is required before better understanding what autonomous weapon systems are and the issues they raise is putting the cart before the horse. Better understanding of the issue is needed before deciding on a path forward.

This paper is intended to frame discussions about autonomy in weapons in order to help inform ongoing debates about the appropriate role of autonomy and human control in the use of force. Too often, arguments about how autonomy ought to be used in future weapons are made without a good understanding of how autonomy is already used today and, in some cases, has been used for decades. In some situations, this can mean individuals advocating for rules governing autonomy in future weapons that would, if implemented, ban wide swathes of existing weapons without any clear evidence that such weapons have been harmful.[54] In fact, in some forms of autonomy, such as those that enable precision-guided weapons, have actually reduced civilian casualties in war.[55] In any case, rules that capture weapon systems that have already been used widely for decades without significant concerns almost certainly miss the mark about what is novel about potential future autonomy in weapons.

Future weapon systems are likely to incorporate greater autonomy in a number of ways, and different types of

54. For example, the International Committee for Robot Arms Control's (ICRAC) proposed "minimum necessary standards" for meaningful human control would, if implemented strictly, ban huge classes of weapons that have been used for centuries dating back to the invention of the catapult. International Committee for Robot Arms Control, "ICRAC statement on technical issues to the 2014 UN CCW Expert Meeting, May 14, 2014, http://icrac.net/2014/05/icrac-statement-on-technical-issues-to-the-un-ccw-expert-meeting/.

55. Michael C. Horowitz and Paul Scharre, "Do Killer Robots Save Lives?" Politico Magazine.

autonomy raise different concerns. Some types of autonomy, such as the ability to maneuver around obstacles or threats, or to deploy evasive maneuvers or non-lethal countermeasures like flares or chaff, raise few if any significant concerns. Others raise potentially profound concerns about the role of humans in the use of force.

Thus, clear terminology that delineates between existing weapon systems that incorporate some autonomy but retain human control over the specific targets being engaged and future autonomous weapon systems that would select and engage targets on their own is an essential precondition for debate. Without a commonly agreed-upon lexicon, those discussing the topic of autonomous weapons may fall victim to the different ways in which the word "autonomy" is used.

It also is worth noting that while this paper has focused on autonomy in engagement decisions for weapon systems, there are many other features of the system that may be equally if not more important. These could be related to autonomy in other, non-engagement related aspects of the system, or not related to autonomy at all.

For example, a static autonomous weapon that could select and engage targets on its own (stationed along a country's border, for instance), would have a significantly different level of risk than a mobile system. And a mobile autonomous weapon system that can only operate in a limited geographic area would have a different level of risk than one that had the ability to patrol over a wide area. Similarly, non-autonomous features could also affect the costs and benefits of utilizing an autonomous weapon. Some environments, such as undersea, are less cluttered and have less potential for civilian casualties. Another important consideration is the target being engaged. Autonomous weapon systems that target objects (i.e. anti-materiel or anti-vehicle autonomous weapons) versus people (i.e. anti-personnel autonomous weapons) also pose meaningfully different levels of risk from the perspective of distinguishing military from civilian targets and avoiding civilian casualties. Finally, the destructive capability of the weapon itself is a significant factor in thinking about risk. A bomber has the capability to cause significantly more destruction in the event of a malfunction or cyber attack than a machine gun. All of these factors are significant and come to bear on how to think about the relative potential costs and benefits of autonomous weapon systems.

How autonomy has been used to date highlights the importance of some of these distinctions. While at least thirty nations operate human-supervised autonomous weapons, they are used in a relatively limited context. They have been used, to date, for static defense of installations or onboard defense of human-occupied vehicles. They target objects (anti-vehicle/materiel), not people (anti-personnel). Perhaps most importantly, human controllers not only supervise their operation in real time but also have physical access to the system so that they can manually disable the system in the event of a malfunction, communications failure or cyber attack.

These issues raise important questions for future research. This paper informs current discussions about autonomous weapons by clarifying terminology and helping focus attention on how potential new uses of autonomy in future weapons could differ from historical uses of autonomy. This should be a launching point for further much-needed discussions among activists, lawyers, ethicists, academics, engineers, and military and defense professionals on the appropriate role of autonomy in future weapon systems.

Appendix A: Definitions of Autonomous Weapons from Selected Reports

The U.S. Department of Defense (DoD) official policy on autonomy in weapon systems, DoD Directive 3000.09, *Autonomy in Weapon Systems*, uses the following definitions:[56]

autonomous weapon system: A weapon system that, once activated, can select and engage targets without further intervention by a human operator.-

This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation.

human-supervised autonomous weapon system: An autonomous weapon system that is designed to provide human operators with the ability to intervene and terminate engagements, including in the event of a weapon system failure, before unacceptable levels of damage occur.

semi-autonomous weapon system: 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.

This includes:

Semi-autonomous weapon systems that employ autonomy for engagement-related functions including, but not limited to, acquiring, tracking, and identifying potential targets; cueing potential targets to human operators; prioritizing selected targets; timing of when to fire; or providing terminal guidance to home in on selected targets, provided that human control is retained over the decision to select individual targets and specific target groups for engagement.

"Fire and forget" or lock-on-after-launch homing munitions that rely on tactics, techniques, and procedures to maximize the probability that the only targets within the seeker's acquisition basket when the seeker activates are those individual targets or specific target groups that have been selected by a human operator.

target selection: The determination that an individual target or a specific group of targets is to be engaged.

Christof Heyns, UN Special Rapporteur on extrajudicial, summary or arbitrary executions, wrote in his April 2013 report:[57]

The measure of autonomy that processors give to robots should be seen as a continuum with significant human involvement on one side, as with UCAVs where there is "a human in the loop", and full autonomy on the other, as with [lethal autonomous robotics (LARs)] where human beings are "out of the loop.

Under the currently envisaged scenario, humans will at least remain part of what may be called the "wider loop": they will programme the ultimate goals into the robotic systems and decide to activate and,

56. United States Department of Defense, Autonomy in Weapon Systems, 3000.09

57. Christof Heyns, "Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions, Christof Heyns," United Nations General Assembly, April 9, 2013, www.ohchr.org/Documents/HRBodies/HRCouncil/RegularSession/Session23/A-HRC-23-47_en.pdf.

if *necessary, deactivate them, while autonomous weapons will translate those goals into tasks and execute them without requiring further human intervention.*

Supervised autonomy means that there is a "human on the loop" (as opposed to "in" or "out"), who monitors and can override the robot's decisions. However, the power to override may in reality be limited because the decision-making processes of robots are often measured in nanoseconds and the informational basis of those decisions may not be practically accessible to the supervisor. In such circumstances humans are de facto out of the loop and the machines thus effectively constitute LARs.

"Autonomous" needs to be distinguished from "automatic" or "automated." Automatic systems, such as household appliances, operate within a structured and predictable environment. Autonomous systems can function in an open environment, under unstructured and dynamic circumstances. As such their actions (like those of humans) may ultimately be unpredictable, especially in situations as chaotic as armed conflict, and even more so when they interact with other autonomous systems.

The terms "autonomy" or "autonomous", as used in the context of robots, can be misleading. They do not mean anything akin to " free will" or "moral agency" as used to describe human decision-making. Moreover, while the relevant technology is developing at an exponential rate, and full autonomy is bound to mean less human involvement in 10 years' time compared to today, sentient robots or strong artificial intelligence are not currently in the picture.

Human Rights Watch, in their 2012 report "Losing Humanity," used the following terminology:[58]

Although experts debate the precise definition, robots are essentially machines that have the power to sense and act based on how they are programmed. They all possess some degree of autonomy, which means the ability of a machine to operate without human supervision. The exact level of autonomy can vary greatly. Robotic weapons, which are unmanned, are often divided into three categories based on the amount of human involvement in their actions:

- *Human-in-the-Loop Weapons*: Robots that can select targets and deliver force only with a human command;
- *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; and
- *Human-out-of-the-Loop Weapons*: Robots that are capable of selecting targets and delivering force without any human input or interaction.

The International Committee of the Red Cross (ICRC), in a report on an expert meeting held in March 2014 on autonomous weapons, explained:[59]

There is no internationally agreed definition of autonomous weapon systems. For the purposes of the meeting, 'autonomous weapon systems' were defined as weapons that can independently select and

58 "Losing Humanity: The Case Against Killer Robots" (Human Rights Watch).
59 "Report of the ICRC Expert Meeting on 'Autonomous weapon systems: technical, military, legal and humanitarian aspects', 26-28 March 2014, Geneva" (International Committee of the Red Cross), May 9, 2014, https://www.icrc.org/eng/assets/files/2014/expert-meeting-autonomous-weapons-icrc-report-2014-05-09.pdf.

Appendix B: Selected Examples of Human-Supervised Autonomous Weapon Systems

SELECTED EXAMPLES OF AIR AND MISSILE DEFENSE SYSTEMS

AEGIS COMBAT SYSTEM

Date of Introduction .... 1983[1]

Manufacturer.............................. Lockheed Martin[2]

Manufacturing Country........... United States

National Operators..................... Australia, Japan, Norway, South Korea, Spain, and the United States[3]

1. United States Navy, "Fact File: Aegis Weapon System," November 22, 2013.
2. Lockheed Martin, "Aegis Combat System."
3. Defense Security Cooperation Agency, "Japan – AEGIS Weapon System Upgrade," December 10, 2012; United States Navy, "Fact File: Aegis Weapon System"; United States Navy Naval Sea Systems Command, "FY 15-19 Aegis Lifetime Support Requirements for Foreign Military Sales," April 30, 2014.

GOALKEEPER

Date of Introduction .... Early 1990s[1]

Manufacturer .... Thales Netherlands[2]

Manufacturing Country .... Netherlands

National Operators .... Belgium, Chile, the Netherlands, Portugal, Qatar, South Korea, the United Arab Emirates, and the United Kingdom[3]

1. Kelvin Fong, "CIWS: The Last-ditch Defence," Asian Defense Journal (July/August 2008), 18-20.
2. Thales Netherlands, "Goalkeeper – Close-in Weapon System."
3. Ibid.

KASHTAN CLOSE-IN WEAPON SYSTEM

Date of Introduction .... 1988[1]

Manufacturer .... KBP Instrument Design Bureau[2]

Manufacturing Country .... Russia

National Operators .... China, India, Russia[3]

1. Kelvin Fong, "CIWS: The Last-ditch Defence," Asian Defense Journal (July/August 2008), 18-20.
2. "Комплекс 3М87 Кортик / Каштан (SA-N-11 GRISON)," Military Russia, January 17, 2009.
3. "Kashtan Close-In Weapon System," Military.com, November 13, 2013.

MK 15 PHALANX CLOSE-IN WEAPON SYSTEM

Date of Introduction .... 1980[1]

Manufacturer .... The Raytheon Company[2]

Manufacturing Country .... United States

National Operators.................... 25 navies including Australia, Bahrain, Canada, Egypt, India, Israel, Japan, New Zealand, Pakistan, Poland, Portugal, Saudi Arabia, South Korea, Taiwan, the United Kingdom, and the United States[3]

1. United States Navy, "MK 15 Phalanx Close-In Weapon System," November 15, 2013.
2. The Raytheon Company, "Phalanx Close-In Weapon System (CIWS)."
3. The Raytheon Company, "Phalanx Close-In Weapon System (CIWS)"; The Raytheon Company, "Raytheon Awarded $31.2 Million from Egypt for Phalanx 1B Ship Self Defense Systems," July 28, 2003; "$45M for Phalanx Upgrades to U.S. Navy, Portugal," Defense Industrial Daily, May 17, 2005; "$169.9M For Support of Phalanx Close-in Ship Defense Guns," Defense Industrial Daily, July 5, 2006; "Raytheon services India's Phalanx," UPI, July 22, 2010; "Raytheon to upgrade US Navy's Phalanx weapon systems," naval-technology.com, September 12, 2013; "Raytheon, South Korea in Phalanx deal," UPI, February 25, 2014.

PATRIOT

Date of Introduction .... 2002[1]

Manufacturer .... The Raytheon Company[2]

Manufacturing Country .... United States

National Operators .... Egypt, Germany, Greece, Israel, Japan, Kuwait, the Netherlands, Saudi Arabia, South Korea, Spain, Taiwan, the United Arab Emirates, and the United States[3]

1. "Patriot Missile Long-Range Air-Defence System, United States of America," army-technology.com.
2. The Raytheon Company, "Patriot."
3. "Patriot Missile Long-Range Air-Defence System, United States of America," army-technology.com; "Fact file: Patriot missile defence," BBC, updated December 4, 2012.

SeaRAM ANTI-SHIP MISSILE DEFENSE SYSTEM

Date of Introduction .... 2008[1]

Manufacturer .... The Raytheon Company[2]

Manufacturing Country .... United States

National Operators .... Japan and the United States[3]

1. "SeaRAM Anti-Ship Missile Defence System, United States of America," naval-technology.com.
2. The Raytheon Company, "SeaRAM Anti-Ship Missile Defense System."
3. "SeaRAM Anti-Ship Missile Defence System, United States of America," naval-technology.com.

SELECTED EXAMPLES OF GROUND VEHICLE ACTIVE PROTECTION SYSTEMS

AMAP-ADS

Date of Introduction .... 2011[1]

Manufacturer .... IBD Deisenroth Engineering[2]

Manufacturing Country .... Germany

National Operators .... "An Asian nation" (rumored to be Singapore)3

1. Rheinmetall Defence, "Rheinmetall takes up a majority share in ADS GmbH," January 2, 2011.
2. IBD Deisenroth Engineering, "AMAPTM-ADS: The Active Protection System."
3. Rheinmetall Defence, "Rheinmetall takes up a majority share in ADS GmbH"; "Singapore's Defence Stance," Asian Military Review, May 1, 2013.

ARENA

Date of Introduction .... 1993[1]

Manufacturer .... KBP Instrument Design Bureau[2]

Manufacturing Country .... Russia

National Operators .... Russia[3]

1. CPT Tom Meyer, U.S. Army, "Active Protection Systems: Impregnable Armor or Simply Enhanced Survivability?," Armor (May/June 1998), 7-11.
2. "ARENA Active Protection System."
3. Meyer, "Active Protection Systems: Impregnable Armor or Simply Enhanced Survivability?"

DROZD

Date of Introduction .... 1983[1]

Manufacturer .... KBP Instrument Design Bureau[2]

Manufacturing Country .... Russia

National Operators .... Russia[3]

1. "Hard-Kill Active Defense Solutions," Defense Update, October 12, 2005.
2. Ibid.
3. Ibid.

DROZD-2

Date of Introduction .... 2005[1]

Manufacturer .... KBP Instrument Design Bureau[2]

Manufacturing Country .... Russia

National Operators .... Russia[3]

1. "Hard-Kill Active Defense Solutions," Defense Update, October 12, 2005.
2. Ibid.
3. Ibid.

IRON CURTAIN

Date of Introduction .... In development[1]

Manufacturer .... Artis, LLC[2]

Manufacturing Country .... United States

National Operators .... United States[3]

1. David Hambling, "New Bomb-Resistant Trucks Will Blast RPGs Before They Hit," WIRED, November 27, 2009.
2. Artis, LLC, "Iron Curtain."
3. Hambling, "New Bomb-Resistant Trucks Will Blast RPGs Before They Hit."

LEDS-150

Date of Introduction .... 2015 (anticipated)[1]

Manufacturer .... Saab[2]

Manufacturing Country .... South Africa

National Operators .... Pakistan and India are reportedly interested in purchasing the LEDS-150[3]

1. Guy Martin, "Saab seeking LEDS customer," defenceWeb, July 22, 2013.
2. aab, "LEDS full spectrum active protection for land vehicles."
3. Martin, "Saab seeking LEDS customer."

QUICK KILL

Date of Introduction .... In development[1]

Manufacturer .... The Raytheon Company[2]

Manufacturing Country .... United States

National Operators .... United States[3]

1. The Raytheon Company, "Raytheon's Quick Kill Active Protection System defeats one of the most lethal armor-piercing Rocket Propelled Grenades," January 9, 2013.
2. The Raytheon Company, "Quick Kill™ Active Protection System (APS)."
3. The Raytheon Company, "Raytheon's Quick Kill Active Protection System defeats one of the most lethal armor-piercing Rocket Propelled Grenades."

SHARK

Date of Introduction .... 2006 (est.)[1]

Manufacturer .... Thales Group and IBD Deisenroth Engineering[2]

Manufacturing Country .... France[3]

National Operators .... France[4]

1. "Active Defense System," Defense Update, July 26, 2006.
2. "Thales, IBD Pursue the SHARK Active Protection System," Defense Update, June 12, 2010.
3. Ibid.
4. Ibid.

TRENCH COAT

Date of Introduction .... 2012 (est.)[1]

Manufacturer .... Rafael Advanced Defensive Systems[2]

Manufacturing Country .... Israel

National Operators .... Israel[3]

1. Elad Benari, "Report: Israel Concerned Hizbullah Has Advanced Anti-Tank Weapon," Arutz Sheva, February 24, 2012.
2. Ibid.
3. Amos Harel, "Precise missiles, cyber-attacks, mass border invasions: Israel's next war," Haaretz, October 12, 2013; Benari, "Report: Israel Concerned Hizbullah Has Advanced Anti-Tank Weapon."

TROPHY

Other Name .... ASPRO-A

Date of Introduction .... 2007[1]

Manufacturer .... Rafael Advanced Defensive Systems[2]

Manufacturing Country .... Israel

National Operators .... Israel[3]

1. "Trophy Active Protection System," Defense Update.
2. Rafael Advanced Defensive Systems, "Trophy."
3. "Trophy Active Protection System," Defense Update.