

# So Just What Is a Killer Robot?

## Detailing the Ongoing Debate around Defining Lethal Autonomous Weapon Systems

AUSTIN WYATT, PHD

Developing a definition for a complete lethal autonomous weapon system (LAWS) is arguably one of the major stumbling blocks to developing an effective international response to the emergence of increasingly autonomous military technology, whether regulation or a developmental ban. As a result of political and practical issues, the international group of experts convened by the United Nations has been unable to generate a definition of autonomous weapon systems that would be universally agreed or operate as the basis for a preemptive development ban. In this gap, various actors from states to arms companies to scholars have developed competing definitions for what they would consider LAWSs.

This article will compare some of these competing definitions, presenting them for consideration of their merits and differences. Whether a given definition would be considered “prominent” in this respect is largely dependent on the extent to which it was cited in the scholarly literature. It would also depend on whether the definition was referred to in the official statements issued after each meeting of the Group of Governmental Experts on LAWSs, and the extent of the author’s broader contribution to military diffusion studies or Autonomous Weapon Systems (AWS) research. This article will draw together elements of competing definitions from scholars, including Ariel Conn, Chris Jenks, and Michael C. Horowitz.<sup>1</sup>

Overall, this article is to present the current state of understanding that underpins the ongoing international debate of AWSs. The core purposes of this article are (1) to present a succinct picture of what AWSs are to demonstrate the importance of differing definitions of this emerging technology; and (2) to present an argument in favor of refocusing the international community on developing objective, commonly held, and function-based understandings of autonomy in the military context.

## **Distinguishing Autonomous Weapon Systems, Unmanned Platforms, and Artificial Intelligence**

Regardless of the specific definition, it is important to note at the outset that it is not realistic to consider autonomy in the robotics field in binary terms; instead, it is much more analytically effective to consider autonomy as a function-based spectrum where human interaction remains present at some point, even if it is limited to the production or strategic deployment stages.<sup>2</sup>

At the time of writing, there have been no publicly acknowledged deployments of fully AWSs. This deficiency is largely due to the ongoing legal and definitional uncertainty. However, a genuine question remains about the feasibility of imbuing a weapon system with capabilities that could be objectively classed as *autonomous*.<sup>3</sup> While there have been deployments of weapon systems that can operate in a manner independent from human supervision (the DoDaam Super aEgis II is an example),<sup>4</sup> a division must be drawn between weapon systems that are truly “autonomous” weapons and those that are merely “highly automated.”<sup>5</sup>

It is also important to note that direct military applications of artificial intelligence and other related technologies comprise only a comparatively minor section of the broader research efforts in these fields. In a reverse of the traditional development burden of an emerging major military innovation, development is primarily occurring outside of the security space. Instead, commercial and university-based research has been principally intended to contribute to civilian projects, such as self-driving cars and home automation. As dual-use technologies, advances in related enabling components are still relevant in outlining our progress toward a future demonstration point of LAWSs. However, in addition to the fact that artificial intelligence software requires task-specific data, military co-option of these technologies would require far more robustness and resistance to interference than is generally present in civilian-designed systems.

### **Definitions of Autonomous Weapon Systems Put Forward by States**

The most common definition of LAWSs originated in a 2012 US Department of Defense (DOD) directive on autonomous weapon systems.<sup>6</sup> This directive outlined the DOD’s view on developing an autonomous capability for weapon systems and the required level of human involvement. This document defines a weapon as fully autonomous if, when activated, it “can select and engage targets without further intervention by a human operator.”<sup>7</sup> Interestingly, DOD Directive 3000.09 lists a requirement for sufficient training for human operators, which indicates a recognition that human operators would have to retain some level of oversight over any use of force decisions. The concern of how to balance the need

to achieve effectiveness in a battlespace characterized by an operational tempo potentially beyond the capacity of human reaction time while also maintaining sufficiently effective human oversight to guard against unintended engagements is apparent in this directive.<sup>8</sup> Finally, DOD Directive 3000.09 also contained a built-in process for obtaining waivers for development, deployment, or even the transfer of LAWSs in situations that potentially contravene the policy.<sup>9</sup> Despite being due to expire at the end of 2017, DOD Directive 3000.09 was still in effect at the time of writing and features prominently in the developing discourse on LAWSs. As the most commonly cited state definition for autonomous weapon systems, the DOD Directive 3000.09 definition has been used as the starting point for the definitions used by multiple other actors, including nongovernmental organizations such as the Campaign to Stop Killer Robots.<sup>10</sup> While this definition has found traction amongst scholars, it has largely been received critically. For example, Heather Roff criticized the DOD definition because the terms *select* and *engage* are open to interpretation.<sup>11</sup> Notwithstanding scholarly critique, the DOD definition is arguably the natural starting point for developing a working definition of AWSs.

Despite its flaws, the DOD definition does represent a more realistic, if non-specific, view of autonomy in weapon systems than the definitions adopted by some other states. In 2011, for example, the UK Ministry of Defence definition referred to autonomous systems having the capability to understand “higher level intent and direction” and that individual actions “may not be” predictable.<sup>12</sup> This definition seems to indicate that a platform or military system must possess artificial intelligence with a level of self-awareness that bleeds into the field of general artificial intelligence (AI). It is highly unlikely that any state actor would countenance the development of weapons that they could not predict, even if it were technologically possible to create LAWSs with the capacity to interpret higher-level intent. The concept of this level of full autonomy has been justifiably dismissed as a distraction in the literature,<sup>13</sup> as an approach driven by this definition simply does not account for the weapon systems that are actually in development.

On 14 April 2018, China became the first permanent member of the Security Council to publicly endorse a ban on the use of LAWSs.<sup>14</sup> This surprise announcement was initially seized on as a victory by the Campaign to Stop Killer Robots and covered extensively in the media, but closer analysis identifies this announcement as an important example of how states can utilize definitional factors to gain influence over the development of LAWSs.

The Chinese definition of LAWSs is based around five characteristics, which serve to exclude other forms of increasingly autonomous military technologies from the discourse. The first characteristic is that a device must carry a “sufficient

payload” and be intended to employ lethal force.<sup>15</sup> While this would obviously cover LAWSs that are designed to directly participate in combat, it would exclude those that carried a less-than-lethal munitions package (such as the remote-operated “Skunkcopter” unmanned aerial vehicle [UAV]), or are designed for an antivehicle/munitions primary function. The second characteristic is an unusually high autonomy barrier, stating that a LAWS would have an “absence of human intervention and control” for the “entire process of executing a task.”<sup>16</sup> China’s statement was vague about what it considers a “task”; this document could refer to a single use of force decision, the acquisition of a target, or an entire deployed mission. Thirdly, and closely linked, the device should have no method of termination once activated to be considered a LAWS.<sup>17</sup> This statement would discount weapon systems that operate autonomously but can be overridden by a human overseer, such as the Phalanx Close-in Weapons System. It is also highly unlikely that a state would deploy a weapon they had no way of deactivating or assuming control over, especially given the comparatively nascent state of AI technology.

The fourth characteristic is that the device must have an indiscriminate effect, that the device would “execute the task of killing and maiming regardless of conditions, scenarios and targets.”<sup>18</sup> This characteristic is an interesting inclusion because international humanitarian law already forbids the use of weapon and weapon platforms that are incapable of being operated in a discriminate manner. The inclusion of this characteristic is complemented by the latter statement in the same announcement that a fully autonomous weapon system would be incapable of satisfying the legal requirement of discriminate use of force. The question of whether a fully autonomous platform could abide international law in the use of discriminate force is central to the debate surrounding LAWSs and has been at the forefront of publicly visible developments in the space. As an example, the Super aEgis II is capable of distinguishing between uniforms and offers clear warnings before engaging to reduce the chances of using lethal force against civilians. Finally, the Chinese definition includes the characteristic that LAWSs could evolve and learn through interaction with the environment they are deployed into in such a way that they “expand its functions and capabilities in a way exceeding human expectations.”<sup>19</sup> This final characteristic leans closer to the UK’s definition of fully autonomous weapons and is effectively arguing that the presence of an actively evolving artificial intelligence is necessary for a weapon system to be considered a LAWS. The concept that LAWSs are being developed with high level AI has been widely criticized by scholars and defense personnel but is a common point raised by concerned nongovernmental organizations (NGO) and smaller states. While it is possible, it is beyond the realm of current technology and

whether states would even be interested in a learning autonomous weapon has been criticized as unrealistic.

There are many reasons that the Chinese definition of *lethal autonomous weapons* is particularly important. Aside from their obvious influence as a permanent member of the security council, autonomous military technology is emerging as a key force multiplier, a factor that is of obvious importance in the context of the Sino-American rivalry and Chinese military modernization. Furthermore, China has a proven track record of using and then ignoring international law as a tactic for advancing its interests, for example, consider China's reaction to being ruled against by the UN permanent court of arbitration in its case against the Philippines over territorial disputes in 2016.<sup>20</sup> Finally, China has already emerged as a major exporter of UAVs (armed and unarmed) to both state and nonstate actors.<sup>21</sup> Indeed, the 2017 decision to reduce export restrictions on US companies was partially motivated by a desire to counterbalance the market dominance achieved by China in the UAV export market. While China's decision to support a ban on the development and use of AWSs seems to be a victory for those opposed to LAWSs, the actual content of their announcement reveals the importance of definitional agreement.

The Chinese announcement clearly excludes large aspects of the developing autonomous military market; however, it has proven quite common in the definitional debate for state and scholarly actors to put forward definitions that have additions that limit the scope of their application. The inclusion of "lethal" in LAWSs excludes weapon platforms that are designed to utilize less-than-lethal ammunition or guide other munitions while the requirement of "higher level" autonomy excludes the plethora of human supervised weapon systems that are already deployed or in development. As encountered by the UN-sponsored Group of Governmental Experts on LAWSs, this disagreement on a common definition hampers efforts to develop either a ban or effective regulatory controls.<sup>22</sup>

Part of the problem is that, while most commonly cited definitions are broadly similar in their top-level language, when one attempts to apply these definitions or questions their underlying assumptions discrepancies emerge. Given the regulatory and discursive power of definitions in this debate,<sup>23</sup> there is a clear political and strategic incentive for states to adopt distinct discursive frames for understanding *autonomy* in this sense. This understanding implies that, among states as a minimum, definition discrepancies are likely to remain,<sup>24</sup> at least while the debate remains focused on the question of a ban.

The complex definitional debate surrounding the term *lethal autonomous weapon system* is one of the key reasons that international efforts to implement a preemptive ban have stalled. Seven states are publicly believed to be developing lethal

autonomous weapon systems: the US, South Korea, China, Russia, India, the United Kingdom, and Israel, though none has admitted to possessing a functioning fully AWS.<sup>25</sup> Only 19 countries publicly support an outright developmental ban; however, this support is based on divergent conceptual understandings of “fully autonomous weapons.” The clear majority of the 63 other states that have publicly stated a position support the continuation of governmental discussions.<sup>26</sup> This support shows that, while the majority of states do not support a preemptive ban, they are concerned and willing to continue high-level discussions toward generating a normative and legal framework to control the impact of LAWSs. Outside the land of government press releases, the 2017 intergovernmental meeting of experts was cancelled, ostensibly due to a lack of funds. The “discussion” advocated by the majority of states in 2019 has therefore been largely organized by NGOs, scholarly communities and regional interstate bodies.

### **Identifying Commonalities in the Focus of Nonstate Definitions of LAWS**

Despite emerging as the principle vehicle for pushing forward discussion on the challenges presented by the emergence of increasingly autonomous weapon systems, there remains definitional disagreement among civil society and scholars, nor has there been any concrete steps taken toward developing an universally agreed set of functional standards for determining whether a given weapon system would fall under the proposed ban.

The majority of actively participating NGOs, including the International Committee for Robot Arms Control, Article 36, Human Rights Watch, and the International Committee of the Red Cross (ICRC), subscribe to functionally similar definitions. This is unsurprising given that these organizations are members of the Campaign to Stop Killer Robots (CSKR), which has been the leading advocate in this space since 2012. Another member of the CSKR—Reaching Critical Will (a Women’s International League for Peace and Freedom program)—defines fully *autonomous weapon systems* as follows:

*“Killer robots are fully autonomous weapon systems. These are weapons that operate without meaningful human control, meaning that the weapon itself can take decisions about where and how it is used; what or whom it is used against; and the effects of use.”*<sup>27</sup>

There are three elements from this definition that can be commonly identified in the published literature and discussion papers produced by NGOs on this issue.

## Lethality

The first element is that this definition explicitly states that fully autonomous weapons are Killer Robots. As part of the campaign's name, this is obviously, a central element of the CSKR's perspective. The term *killer robot* is a dysphemism that has been consistently used to focus the discourse on the capability of lethal aspect of LAWSs, particularly in media appearances and published materials, as well as in the central questions of the public surveys commissioned by CSKR over the past three years. While a legitimate and important concern, the lethal use of AWSs is the most controversial potential use of the underlying technologies and arguably distracts from the rapid progress that states are making on systems that are not designed primarily for the use of lethal force. Ajey Lele has argued that focusing on lethality makes it impossible to come to a "foolproof" definition because sometimes the lethality of an autonomous system will depend on the purpose of its deployment.<sup>28</sup> Heather Harrison Dinniss also argued that the purpose of deployment, target justification, and user intention were more important than the weapon's inherent nature.<sup>29</sup>

While Lele referred specifically to cyberwarfare, other problematic autonomous systems could include AI-enabled battlefield decision-making aides, cyber warfare agents, and "support" unmanned ground vehicles whose stated purpose is for battlefield resupply, none of which would necessarily be covered by a ban that followed this definition, yet could be used in a manner that leads to death and injury.

### "Full" Autonomy and Critical Functions

Secondly, it is problematic to focus on whether a hypothetical system having full autonomy. While distinguishing fully autonomous systems from platforms that clearly operate under human supervision or within functional constraints has clear utility (at least from a policymaking perspective), autonomy is not a binary characteristic that can be easily identified, separated and measured. Jenks argues that it is more effective to consider autonomy as the "capability of the larger system enabled by the integration of human and machine abilities" and that autonomy (even in weapon systems) is inherently bounded by the interaction between human and machine.<sup>30</sup> Alternatively, Horowitz has argued that AI (the most important underlying technology for autonomous systems) is better conceptualized as a disruptive enabling technology rather than a distinct weapon system, maintaining that AI is conceptually closer to the combustion engine than the aircraft carrier.<sup>31</sup>

It is therefore important to focus on the extent to which a system has control over its *critical functions* independent of human intervention or supervision, which is reflected in the *Reaching Critical Will* definition. The critical functions of a weapon system are the processes used to select, acquire, track and attack targets.<sup>32</sup> These processes are considered critical because they become the core of the kill chain once human supervision is removed.<sup>33</sup> The *kill chain* is a commonly used term within the US military and in the relevant academic literature. The level of control over these functions is central to the ICRC definition of autonomous weapon systems.<sup>34</sup> Similarly for Anderson, it is the capacity of autonomous weapons to “undertake” the process of identification, rather than merely to respond to a particular stimulus that is their primary characteristic.<sup>35</sup> By focusing on the *critical functions* of the weapon system, advocates of a ban took a step toward the functional benchmarks that would be required for effective international regulation of LAWSs.

### **Meaningful Human Control**

The final commonly seen element that can be extracted from the *Reaching Critical Will* definition is the importance placed on retaining a *Meaningful Human Control* standard. The concept of *Meaningful Human Control* arose as a response to the perceived “accountability gap” with autonomous weapon systems and has been a major talking point at each meeting of experts.<sup>36</sup> The Campaign to Stop Killer Robots, and affiliated groups, have enthusiastically embraced Meaningful Human Control as a vital standard that, employed alongside a ban on fully autonomous weapons, would arguably prevent the transfer of the decision to use lethal force to those robotic systems that are not prohibited. However, despite this prominence, there remains no universal agreement on the limits of its meaning or how to ensure that it is maintained. For example, Christof Heyns has written that autonomous law enforcement weapons would still be under meaningful human control if a human authorised that specific target and instance of force, even if the weapons did not engage immediately.<sup>37</sup> The literature has begun to push back against this lack of definitional clarity, as well as the murkiness surrounding definitions of autonomy in the military context.<sup>38</sup> As a prominent example, R. Crotoft has challenged the blind acceptance of Meaningful Human Control.<sup>39</sup> Instead, her work explores how the concept of Meaningful Human Control would interact with inconsistent domestic state laws as well as international humanitarian law.<sup>40</sup>



Furthermore, tragic historical examples with semiautonomous weapon systems, including the downing of Iran Air Flight 655, demonstrated that Meaningful Human Control must be paired with robust verification procedures and organizational modifications, including comprehensive operator and commander training. Without these measures, there is a danger that human supervisors would operate on the basis of overly enthusiastic interpretations of the platform's capability, even where "meaningful human control" is theoretically maintained.<sup>41</sup>

### **So, What is an Autonomous Weapon System?**

Attempting to present an authoritative single definition of LAWSs in the midst of the ongoing international debate would be a hubristic goal for this article. As with terrorism, the broad strokes of a definition have been admirably outlined by others and are generally agreed, the continued international debate centers on the specifics and is sustained by discursive differences that are primarily political in nature. However, by drawing on the positions explained above, and a selection of definitions developed by prominent scholars, it is possible to synthesize a working definition that would be sufficient to facilitate discussion separate from the politicized CCW process.

At its most simplistic an AWS could be thought of as a computer that is analyzing data inputted from multiple conventional sensors to inform its actions without direct human involvement. While insufficiently detailed, this kind of definition is useful for scholars whose analysis is focused on the ethical, moral, strategic, or legal issues raised by LAWSs. For example, Maya Brehm adopted a basic definition of AWSs as "a weapon system with sensors, algorithms and effectors," with the explicit acknowledgement that this approach sidestepped the ongoing debate while providing a sufficient descriptive picture for the reader. However, for regulation to be effective, it would require a more operationalizable and detailed approach.

At the core of this approach should be a consideration of the level of independent control that a system exercises over its critical functions.<sup>42</sup> Setting aside those weapon systems that are either inert (requiring human operation) or automated (such as landmines),<sup>43</sup> this approach would help identify whether a system is operationally semiautonomous, supervised by a human operator, or exercises operationally full autonomy over its critical functions. Interestingly, existing definitions have placed emphasis on different critical functions in their approach to autonomous weapon systems. For example, Crotoft emphasized

the weapon's ability to process information to make targeting decisions,<sup>44</sup> while Horowitz emphasized the ability to select a target that was not preselected by an operator.<sup>45</sup>

Furthermore, given the goal is to create a definition suitable for the development of technical standards among states that are currently pursuing AWSs, as well as potential future importers, it is better to focus the definition on autonomy at the platform level, rather than disposable munitions or systems where autonomous agents completely replace humans in the planning of military action.<sup>46</sup>

Based on these features, consider the following as an early example of such a working definition for LAWSs:

“A fully autonomous Lethal Autonomous Weapon System (LAWS) is a weapon delivery platform that is able to independently analyze its environment and make an active decision whether to fire without human supervision or guidance.”<sup>47</sup>

This is just one definition for your review. As Jenks noted, shifting international discussion away from calls for and against a preemptive ban toward discussing measurable technical standards around these critical functions is unlikely to resolve the currently stalled process at the CCW.<sup>48</sup> However, it is now approaching seven years into the CCW meeting of experts discussions without any agreement on objective standards for measuring autonomy or even a definition around which regulation could be meaningfully discussed. The development of autonomous military technology has not comparably slowed during this process, bringing us closer to the introduction of fully autonomous military technology without a common definitional basis from which a governing framework could be effectively developed.

## **Conclusion**

In conclusion, the continued debate surrounding the challenge of defining AWSs highlights the need to reconsider the international community's current approach. Instead, the scholarly community should refocus on exploring and defining the line that must be drawn between true functional autonomy and mere sophisticated automation within the current paradigm of conflict. This line can sometimes be difficult for policymakers, academics, and even practitioners to see; however, it is vital that we distinguish such systems, as

well as the use of AI-enabled technologies for logistics purposes from the term *lethal autonomous weapon system*. As Horowitz has argued elsewhere,<sup>49</sup> definitions have power and political significance. The international community cannot continue to focus debate solely around the question of a ban under international law. Instead, scholars and policymakers need to take a step back and take the time to develop objective, replicable, and agreeable standards for determining whether a weapon system is autonomous, merely automated, or falls into a different category.

If the current trajectory continues, the general understanding of autonomous weapon systems risks following the example of terrorism, which is still lacking a universal definition almost 20 years after 9/11. Without this agreement, any international regulation would be vulnerable from its inception. If this discussion cannot effectively take place in the forum of the United Nations, whether due to continued resistance from certain states or otherwise; it is time that regional security organizations step up to meet this gap. A concrete, function-based definition, agreed to between regional middle power states, would be an applaudable first step, and perhaps regional organizations could lead the way.<sup>50</sup>

**Austin Wyatt, PhD**

Dr. Wyatt (PhD, Australian Catholic University) is a research associate in the Values in Defence and Security Technology group at The University of New South Wales at the Australian Defence Force Academy. His research concerns autonomous weapons with a particular emphasis on their disruptive effects in Southeast Asia.

## Notes

1. Ariel Conn, “The Problem of Defining Autonomous Weapons,” Future of Life Institute, 30 November 2016, <https://futureoflife.org/>; Chris Jenks, “The Distraction of Full Autonomy & the Need to Refocus the CCW Laws Discussion on Critical Functions,” in *Legal Studies Research Paper* (SMU Dedman School of Law, 2016); and Michael C. Horowitz, “Why Words Matter: The Real World Consequences of Defining Autonomous Weapons Systems,” *Temp. Int’l & Comp.* 30, 2016.

2. Kenneth Anderson, “Why the Hurry to Regulate Autonomous Weapon Systems—But Not Cyber-Weapons,” *Temple International and Comparative Law Journal* 30, no. 1 (2016): 17–42.

3. Anderson, “Why the Hurry to Regulate Autonomous Weapon Systems.”

4. Simon Parkin, “Killer Robots: The Soldiers That Never Sleep,” BBC News, 2015, <http://www.bbc.com/>.

5. Anderson, “Why the Hurry to Regulate Autonomous Weapon Systems,” 17–42.

6. Department of Defense (DOD), Directive 3000.09, 21 November 2012, <https://www.esd.whs.mil/>.

7. DOD Directive 3000.09.

8. DOD Directive 3000.09 defines “unintended engagements” as “the use of force resulting in damage to persons or objects that human operators did not intend to be the targets of U.S. military operations.” DOD Directive 3000.09.

9. DOD Directive 3000.09.

10. Campaign to Stop Killer Robots, “The Threat of Fully Autonomous Weapons,” accessed 26 May 2020, <https://www.stopkillerrobots.org/>.

11. Heather Roff quoted in Conn, “The Problem of Defining Autonomous Weapons.”

12. UK Ministry of Defence, “Joint Doctrine Note 2/11: The UK Approach to Unmanned Aircraft Systems,” ed. UK Ministry of Defence, 2011.

13. Jenks, “The Distraction of Full Autonomy.”

14. Elsa B. Kania, “China’s Strategic Ambiguity and Shifting Approach to Lethal Autonomous Weapons Systems,” *Lawfare*, 17 April 2018, <https://www.lawfareblog.com/chinas-strategic-ambiguity-and-shifting-approach-lethal-autonomous-weapons-systems>.

15. Chinese Delegation to CCW, “Position Paper,” 2018.

16. Chinese Delegation to CCW, “Position Paper.”

17. Chinese Delegation to CCW, “Position Paper.”

18. Chinese Delegation to CCW, “Position Paper.”

19. Chinese Delegation to CCW, “Position Paper.”

20. Laura Zhou, “China’s Foreign Ministry Joins War of Words against Singapore over South China Sea Dispute,” *South China Morning post*, 27 September 2016.

21. Elisa Catalano Ewers et al., “Drone Proliferation: Policy Choices for the Trump Administration,” in *Papers for the President* (Center for a New American Security, 2017).

22. Human Rights Watch, “‘Killer Robots’: Russia, US Oppose Treaty Negotiations: New Law Needed to Retain Meaningful Human Control over the Use of Force,” 19 August 2019, <https://www.hrw.org/>.

23. Horowitz, “Why Words Matter,” 85.

24. Maya Brehm, "Defending the Boundary: Constraints and Requirements on the Use of Autonomous Weapon Systems under International Humanitarian and Human Rights Law," Geneva Academy Briefing 9, 2017, <https://www.geneva-academy.ch/>.

25. International Committee of the Red Cross (ICRC), "Views of the International Committee of the Red Cross (ICRC) on Autonomous Weapon Systems," (paper presented at the Convention on Certain Conventional Weapons (CCW) Meeting of Experts on Lethal Autonomous Weapons Systems (LAWS), Geneva Academy, 11–15 April 2016.

26. Campaign to Stop Killer Robots, "Country Views on Killer Robots," 11 October 2017, news release, 11 October 2017, <http://www.stopkillerrobots.org/>.

27. Ray Acheson, "A WILPF Guide to Killer Robots," Reaching Critical Will, 2019.

28. Ajey Lele, "Debating Lethal Autonomous Weapon Systems," *Journal of Defence Studies* 13, no. 1 (2019): 51–70, <https://idsa.in/jds/>.

29. Heather Harrison Dinness, *Cyber Warfare and the Laws of War* (Cambridge: Cambridge University Press, 2012), quoted in Brehm, "Defending the Boundary: Constraints and Requirements on the Use of Autonomous Weapon Systems under International Humanitarian and Human Rights Law."

30. Jenks, "The Distraction of Full Autonomy."

31. Michael C. Horowitz, "Artificial Intelligence, International Competition, and the Balance of Power," *Texas National Security Review* 1, no. 3 (2018).

32. ICRC, "Autonomous Weapon Systems: Technical, Military, Legal, and Humanitarian Aspects," in *Expert Meeting*, 2014.

33. Curtis E. LeMay Center for Doctrine Development and Education, "Annex 3–60 Targeting," ed. Air Education and Training Command (AETC) (Maxwell AFB, Alabama: AETC, 2017); and Julian C. Cheater, *Accelerating the Kill Chain Via Future Unmanned Aircraft*, ed. Centre for Strategy and Technology, Air War College, Blue Horizons Paper, 2000.

34. "Any weapon system with autonomy in its critical functions. That is, a weapon system that can select (i.e., search for or detect, identify, track, select) and attack (i.e., use force against, neutralize, damage, or destroy) targets without human intervention." "International Humanitarian Law and the Challenges of Contemporary Armed Conflicts," International Committee of the Red Cross, 2015.

35. Anderson, "Why the Hurry to Regulate Autonomous Weapon Systems."

36. A. Krishnan, "Automating War: The Need for Regulation," *Contemporary Security Policy* 30, no. 1 (2009): 172–93, <https://www.tandfonline.com/>; Benjamin Kastan, "Autonomous Weapons Systems: A Coming Legal 'Singularity'?" *Journal of Law, Technology & Policy* 45, no. 1, 2013, 45–82, <http://illinoisjltp.com/>; Daniel N. Hammond, "Autonomous Weapons and the Problem of State Accountability," *Chicago Journal of International Law* 15 no. 2, 2014, <https://chicagounbound.uchicago.edu/>; Markus Wagner, "The Dehumanization of International Humanitarian Law: Legal, Ethical, and Political Implications of Autonomous Weapons Systems," *Vand. J. Transnat'l Law* 47, no. 1371, 2014, <https://papers.ssrn.com/>; and James I. Walsh, "Political Accountability and Autonomous Weapons," *Research & Politics* 2, no. 4, 2015, <https://journals.sagepub.com/>.

37. Christof Heyns, "Human Rights and the use of Autonomous Weapons Systems (AWS) During Domestic Law Enforcement," *Human Rights Quarterly* 38, (2016): 350–378.

38. Anderson, "Why the Hurry to Regulate Autonomous Weapon Systems."
39. R. Crootof, "A Meaningful Floor for 'Meaningful Human Control,'" *Temple International & Competitive Law Journal* 30.
40. R Crootof, "A Meaningful Floor."
41. Crootof, "A Meaningful Floor."
42. Lele, "Debating Lethal Autonomous Weapon Systems."
43. Rebecca Crootof, "Autonomous Weapon Systems and the Limits of Analogy." *Harvard National Security Journal* 9 (2018): 51–83, <https://harvardnsj.org/>.
44. "A weapon system that, based on conclusions derived from gathered information and pre-programmed constraints, is capable of independently selecting and engaging targets." Rebecca Crootof quoted in Horowitz, "Why Words Matter."
45. Michael C. Horowitz, "The Ethics & Morality of Robotic Warfare: Assessing the Debate over Autonomous Weapons," *Daedalus* 145, no. 4 (2016).
46. Horowitz, "Why Words Matter."
47. "A fully autonomous Lethal Autonomous Weapon System (LAWS) is a weapon delivery platform that is able to independently analyse its environment and make an active decision whether to fire without human supervision or guidance." Austin Wyatt and Jai Galliot, "Closing the Capability Gap: ASEAN Military Modernization during the Dawn of Autonomous Weapon Systems," *Asian Security* 16, no. 1, 26 September 2018, 53–72, <https://www.tandfonline.com/>.
48. Jenks, "The Distraction of Full Autonomy."
49. Horowitz, "Why Words Matter."
50. Wyatt and Galliot, "Closing the Capability Gap."