The War on “Killer Robots”: Goals and Strategies in the Transnational Campaign Against the Offensive Use of Lethal Autonomous Weapon Systems

John Fan

Swarthmore College, zfan1@swarthmore.edu

Follow this and additional works at: <http://works.swarthmore.edu/swarthmoreirjournal/>

Recommended Citation

Fan, John. “The War on ‘Killer Robots’: Goals and Strategies in the Transnational Campaign Against the Offensive Use of Lethal Autonomous Weapon Systems.” *Swarthmore International Relations Journal* 2 (2017): 50-63.

Available at: <http://works.swarthmore.edu/swarthmoreirjournal/vol1/iss2/9>

This work is brought to you for free and open access by Works. It has been accepted for inclusion in *Swarthmore International Relations Journal* at Swarthmore College by an authorized administrator or Works. For more information, please contact myworks@swarthmore.edu

The War on “Killer Robots”: Goals and Strategies in the Transnational Campaign Against the Offensive Use of Lethal Autonomous Weapon Systems

John Fan

Swarthmore College

At the opening of the 2015 International Joint Conference of Artificial Intelligence, AI scholars presented an open letter, calling for a preemptive ban on autonomous robots that can choose and kill targets without human intervention (Future of Life Institute 2015). The letter cautions against a “global AI arms race” that will eventually make killing easier for terrorists, dictators, and warlords—once autonomous weapons can be mass produced as the technology matures. Even though the kind of fully autonomous weapons system described in the letter has not been developed yet, the letter warns that the development and deployment of lethal autonomous weapons systems (LAWS) are possible in the next few years. The letter was signed by leading AI researchers such as the CEO of Google DeepMind Demis Hassabis, as well as famous academics and technology experts such as physicist Stephen Hawking, CEO of SpaceX Elon Musk, Apple co-founder Steve Wozniak, and linguist, philosopher, cognitive scientist and social critic Noam Chomsky. As of November 2016, the letter has been signed by more than 20,000 people (Ibid.).

Even though many scholars and technology experts have voiced concerns about the development and future deployment of LAWS, discussion of a preventative ban in the UN has been largely unfruitful (Vilmer 2016). This paper views the inability to reach international consensus on a LAWS ban from the perspective of “transnational activists” (Tarrow 2005; Keck and Sikkink 1998), and attempts to answer the following two questions. First, what effects on domestic policy and international law should activists strive for? Second, what available strategies can transnational activists adopt to address the political impasse on the international arena?

On the first question, this paper argues that activists should push for a “ban + regulation” framework that incorporates both a partial ban on LAWS and a set of regulations that mandate their responsible use. The necessity and appropriateness of the framework are motivated by ethical and legal objections to the use of LAWS without direct human control. On the second question, this paper embraces a “two-tier approach” of transnational advocacy that addresses various causes of the impasse from both the domestic level and the international level.

This paper uses the definition of LAWS proposed by Heather Roff: LAWS are *learning* machines that can *autonomously target and fire* without human intervention (Roff 2014, 212-214). Although very broad, a more precise definition is unnecessary. As this paper shall demonstrate later, disagreement on the definition of LAWS has been a significant barrier to international consensus on the regulation of LAWS, and the international community has come to recognize that a comprehensive definition is too early at this point (Vilmer 2016). The

definition that this paper adopts should be sufficient for laying out the groundwork of LAWS policy discussion. The following sections explore the problems that transnational activists are faced with and propose some tentative solutions to the identified problems.

The Problems: The History and Contemporary Politics of LAWS

The use of military “robots” in warfare traces back to World War I, when the United States designed a gyroscope-guided missile called the “Bug,” and Germany designed remote-controlled motorboats that could be fit with explosives (McCormick 2014). However, military AI research really began when the Pentagon gave Massachusetts Institute of Technology significant funding in 1963 (Ibid.). Much progress has been made worldwide since. By 2001, the United States had fully developed drones that could carry Hellfire missiles (Ibid.). In 2005, due to legal and technical concerns, the United States military canceled the plan to build a cruise missile that could autonomously kill targets on a battlefield, even though the technology for such a weapon had long existed (Gubrud 2015). In 2006, South Korea unveiled sentry robots that could automatically track and engage targets, although it was reported that human approval was required before the robots could fire (McCormick 2014).

During the past several years, artificial intelligence research has achieved results that were almost unimaginable in the past. *The Economist* (2016) reports that, due to the use of new techniques such as “deep learning” and “neural networking,” the year 2012 marked the beginning of several groundbreaking developments in the field of artificial intelligence. In 2012, a research team at the University of Toronto significantly improved the image identification ability of previous AI algorithms. Three years later in 2015, an AI algorithm beat the average human level of 95% accuracy in image identification for the first time in history. The same report also mentioned that in 2016, DeepMind, the AI research branch of Google’s parent company Alphabet, made headlines when its AlphaGo system beat Lee Sedol, one of the best human players in the world, in the ancient Chinese board game Go (Ibid.). For human players, the game of Go, in comparison to chess, is a much more intuitive rather than calculative game. Even though chess has already been “solved” in the past by machines using brute force algorithms, solving the game of Go was considered “a grand challenge for AI” that would lead to significant development of AI technology in many areas of application (Gelly et al. 2012, 107). These new techniques of artificial intelligence will most likely find military applications and accelerate the development of LAWS.

The quick development of AI technology supports the researchers’ conclusion in their open letter that “artificial intelligence (AI) technology has reached a point where the deployment of such systems is — practically if not legally — feasible within years, not decades” (Future of Life Institute 2015). Military AI and the deployment of LAWS will revolutionize robotic warfare. In response to this trend, different institutions and organizations have reacted differently. For example, the Obama administration issued a guideline on the responsible use of autonomous weapons systems (Department of Defense 2012), even though the directive is described by one commentator as “unremarkable in substance and arguably should apply to any weapons system” (Gubrud 2015). Its definition of LAWS has also been criticized for lack of

clarity (Roff, quoted in Conn 2016). UN Special Rapporteur Christof Heyns submitted a report on LAWS to the Human Rights Council in 2013, and the report's main recommendations included national moratoria on the development of LAWS and a panel study of the technology's implications (United Nations Human Rights Council, 21). Since 2014, during the annual UN Convention on Certain Conventional Weapons (CCW), experts and countries have met and discussed disarmament of LAWS (Vilmer 2016). However, very little consensus has been reached during the three meetings, due to technical and political reasons (Sayler 2016). First, since most of the discussion revolves around whether LAWS should be banned or not, states have been cautious with their definitions, and little progress has been made as a result (Ibid.). Second, despite a shared understanding that LAWS can be dangerous, many states are still unwilling to adopt a binding framework or implement a preventative ban just yet (Vilmer 2016; Sayler 2016).

The “Accountability Problem” and the “Strategic Robot Problem”

Without a robust regulatory framework or a comprehensive ban, the deployment of LAWS can significantly challenge the force of international law and undermine the military command and control structure. LAWS can cause what can be termed the “accountability problem”: assigning blame (legal or moral) is impossible or extremely difficult after an accident happens. Unless the artificial intelligence that underlies an autonomous weapons system is capable of “higher-order intentionality” about its own beliefs, it usually cannot be morally (or legally) responsible for the harm it causes (Dennett 1996, 354). The question, then, is who else should be responsible for the harm that LAWS may cause, especially since personal accountability is central to international law (Human Rights Watch 2015, 13).

Unless a set of procedures is in place that ensures human involvement whenever an important tactical or strategic decision is made, meaningful personal accountability is nearly impossible. In particular, if human intervention is impossible after LAWS is deployed,¹ and the system violates the international law by, for example, not respecting the humanitarian law requirement of proportionality, either the person(s) who chose to deploy the weapon (the commander), or the person(s) who programmed the system (the programmer) should take responsibility. Lewis (2015, 1324) concludes that the first option is appropriate based on a similar requirement in landmine regulation. However, unless the commander was fully aware that the deployment of LAWS will or will likely lead to violation of international law, it seems unreasonable that the commander should be blamed for their unpredictability, especially if the use of LAWS is permitted by international law. As for landmines, there is already a well-

¹ Note that even semi-autonomous weapons systems defined in the DOD directive does not require human intervention after deployed, and semi-autonomous weapons systems are allowed to use lethal force under the current guideline. See Department of Defense, *Department of Defense Directive 3000.09: Autonomy in Weapon Systems* (Arlington, VA: United States Department of Defense, 2012), 3, <https://www.hsdl.org/?abstract&did=726163>.

established norm against their use, and therefore it is appropriate to blame the commander for any unjustified harm that his or her deployed landmines cause.

The second option also seems inappropriate, since LAWS are unlike other automated weapons that “respond to a preprogrammed set of constraints” (Roff 2014, 212). The programmer cannot fully control the autonomous learning process of LAWS. The programmer need not even be familiar with the international law, just as the developers of AlphaGo did not need to know how to play Go (Gibney 2016, 445).² In conclusion, unless human control is present in every step of LAWS operation, assigning criminal responsibility is almost impossible, and international humanitarian law violations may not be punishable.

Even if the “accountability problem” can be overcome, LAWS that are used without a high level of human involvement can still create what Roff calls the “Strategic Robot Problem” (Roff 2014). Since the battlefield evolves very rapidly and the list of targets must be constantly reviewed and updated by commanders, the deployed LAWS must either possess preprogrammed lists of targets, or be “commanders” that can generate their own target lists. In the first case, LAWS must frequently be recalled to receive updated lists, undermining the claim that LAWS can be cost effective. In the second case, LAWS might generate conflicting military goals with each other, and the multiplicity of strategic actors can undermine the command and control structure and make interoperability, an important element of modern warfare, “mere fiction” (Ibid., 219-220). Moreover, since LAWS operate in isolation and cannot be held morally responsible for their decisions, “moral authority and responsibility [...] vanishes” (Ibid., 220). For those who believe that moral and criminal accountability are normatively indispensable, both the “accountability problem” and the “strategic robot problem” weigh heavily against the use of LAWS without a high level of human control and a robust command and control structure.³

The Solutions: The “Ban + Regulation” Model

Many scholars and experts have advocated for a ban on LAWS (Future of Life Institute 2015). However, as the past three meetings at CCW show, the disagreement over definition and the lack of political will to comprehensively ban LAWS have resulted in a gridlock (Vilmer 2016; Sayler 2016). Arguing that regulation, rather than a ban, is more effective in ensuring compliance and protecting human life, John Lewis (2015, 1310) advocates for LAWS regulation modeled after that of landmine use. However, dividing the available action space into ban versus regulation presents a false dichotomy. This section argues that a “ban + regulation” model is more politically practicable and directly deals with the technical problems mentioned in the previous section. Roff (2016, 123) also concedes that regulation is important if a ban cannot be

² In the case of AlphaGo, it was a general-purpose algorithm and was not preprogrammed with any Go paradigms.

³ For more discussion on the ethics of LAWS, see 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 (2012), 687-709, doi: 10.1017/S1816383112000768.

accomplished, and that we should think about how ban and regulation can together eliminate the dangers of LAWS. This section offers some thoughts on this suggestion.

First, due to the nature of AI technology, a “comprehensive” ban on LAWS without regulatory support is difficult, if not impossible. As Owen (2016) argues, technology such as artificial intelligence cannot be “banned” in the usual sense, especially since AI is a “dual-use” technology that has already been developed for peaceful use. He argues that procedural regulations can be more effective, since it is the individuals who are accountable (Ibid.). Moreover, in contrast to its physical manifestations as “devices,” technology is fluid and cannot be neatly “boxed in” by a comprehensive ban. For example, Marshall (1997, 1392) observes that even though antipersonnel laser weapons that purposefully blind or severely impair soldiers’ vision have been banned by international treaty law, many other military devices that use the technology but are not explicitly “antipersonnel” are left out from the treaty and can still blind soldiers. Examples include “range finders, target illuminators, and anti-sensor systems.” Since it is the devices rather than the technology that are prohibited, the “ban” on laser weapons cannot completely eliminate the inherent risk of laser technology.

Lewis (2015, 1323-1324) argues that LAWS regulations should follow the example of those of landmines. In his view, regulations should take into account their technical details and capacity, appropriate environment for their deployment, possible evasion techniques and mistargeting, level of human control, and other factors. Borrowing from landmine regulations is helpful, especially since there is a high level of compliance with landmine regulations among signatories (Bryden 2013, cited in Lewis 2015, 1318). On the issue of human control, Roff and Moyes (2016) prefer the concept of “meaningful human control” that ensures informed human control and possibility for intervention in every step of the process.

In addition to a regulatory framework, a partial ban on certain possible kinds of LAWS that are clearly dangerous or unacceptable should still be implemented. Lewis (2015) argues against banning LAWS, citing again the landmine example where many states were unwilling to accept a ban on landmines and instead opted for moderate regulations of the Amended Protocol II (1318-1319). However, Article 3 of the Amended Protocol II, in fact, prohibited the use of non-detectable landmines, self-deactivating landmines or landmines that are designed to “cause superfluous injury or unnecessary suffering (United Nations 1996, 135-136). Following this example, LAWS that cause “unnecessary suffering” as part of their designs should also be banned.

A Two-Tier Approach

How, then, should activists strive to achieve an international regulatory regime based on the “ban + regulation” model? This section of the paper draws from the constructivism literature and transnational advocacy network (TAN) literature in IR theory, and argues that activists should adopt a more aggressive and creative approach of norm entrepreneurship and international contention.

As constructivist theorist Alexander Wendt (1992, 397) argues, relationships between states are characterized by “intersubjective understandings and expectations” rather than just

pure material interests. Kenneth Waltz's conception of "self-help" (1979), according to Wendt, is only one of many different ways the international system could be organized (Wendt 1992, 400). Wendt further argues that "positive interdependence of outcome" can create new understandings and expectations in the form of social norms, and commitment to these norms can supersede egoism in states' behavior (Ibid., 417). If constructivism correctly describes the ontology of international relations, then establishing a norm of LAWS disarmament may prevent a global LAWS arms race and lead to more responsible military use of AI technology.

Neo-realists who believe world politics is fundamentally a system of self-help, however, would criticize this vision as hopelessly utopian. As Waltz (1979, 102-105) argues, the international system encourages power-seeking behavior and pursuit of relative gain. Both characteristics tend to encourage a global arms race of LAWS, especially since autonomous weapons have been regarded as "the third revolution in warfare" (Future of Life Institute 2015). Even Wendt does not "contest the neorealist description of the contemporary state system as a competitive, self-help world" (1992, 396), and cautions that the transformation of interest and identity faces numerous constraints (Ibid., 418). However, the example of nuclear nonproliferation shows that pessimism or even fatalism is unwarranted. As Sagan (1996/97, 71-73, 82-86) observes, domestic actors and constraining international norms can both lead to nonproliferation. There is also a strong international norm against the use of nuclear weapons (Tannewald 1999, 435). Moreover, between the United States and Russia, significant progress has been made in nuclear arms control (Arms Control Association 2014). Thus, self-constraining behavior is possible if there is an established norm of disarmament or nonproliferation that stems from a shared understanding of the weapon's danger.

Jean-Baptiste Jeangène Vilmer's (2016) report of the third and most recent meeting on LAWS at the CCW shows a shared understanding among the states of the potential risk of LAWS, but the meeting fell short of building a consensus on its preventative ban. Russia, in particular, only wanted a "discussion" on LAWS at this point rather than negotiation of a formal framework, and even states that opposed LAWS lacked the political will to advocate for a moratorium or preventative ban (Ibid.). In other words, beyond the shared understanding that LAWS can be dangerous, states did not feel obligated to agree to a ban. Furthermore, even without a ban, there was significant disagreement on what types of regulations should apply (Ibid.). For activists who want to create an international norm against the use of LAWS, a two-tier approach that focuses both on the domestic level and the international level is appropriate.

On the domestic level, activists should try to establish a norm against the use of LAWS. There have been numerous works published on "norm entrepreneurship" by constructivist and transnational advocacy networks (TAN) scholars. Norms are defined as "standard[s] of appropriate behavior for actors with a given identity," such as the norm against the use of landmines (Finnemore and Sikkink 1998, 891). According to Martha Finnemore and Kathryn Sikkink (Ibid., 893), "[m]any international norms began as domestic norms and become international through the efforts of entrepreneurs of various kinds." Before norms are institutionalized as international rules and organizations, norm entrepreneurs must dramatize or even "create" issues by using language that resonates with preexisting moral and cultural

understandings, a process usually referred to as “framing” (Ibid., 897-901; Tarrow 2005, 61; Keck and Sikkink 1998, 19). Successful framing should depict the issue in unequivocal terms, and make experts’ technical information easily digestible for the targeted publics and policymakers (Keck and Sikkink 1998, 19). Successful framing should also attempt to “graft” the new moral standards onto preexisting ones, especially in disarmament politics (Carpenter 2007, 103-104). Moreover, since a significant role of transnational activism is to provide information and argue for a moral position, a useful frame must “show that a given state of affairs is neither natural nor accidental, identify the responsible party or parties, and propose credible solutions” (Keck and Sikkink 1998, 19).

Figure 1. A webpage on the website of a campaign against LAWS. *Source:* Campaign to Stop Killer Robots, “Learn,” *Campaign to Stop Killer Robots*, 2015, <https://www.stopkillerrobots.org/learn/>.

Figure 2. A webpage on the the website of a campaign against LAWS. *Source:* Campaign to Stop Killer Robots, “The Problem,” *Campaign to Stop Killer Robots*, 2015, <https://www.stopkillerrobots.org/the-problem/>.

What Similarities Do Killer Robots Share with the Terminator?

- Both can select and fire targets on their own, without human intervention.
- Both are based on artificial intelligence for decision-making.
- Both are machines that cannot be held morally responsible for their killings.
- Neither can be bargained or reasoned with, or feel pity, remorse, fear, or stop before they kill the target.

Do you know?

- Under the Department of Defense guideline, the Terminator is classified as “semi-autonomous” and is allowed to use lethal force.

Figure 3. Using cultural symbols such as the Terminator to illustrate the problem of LAWS.

Bearing in mind these lessons, this paper proposes that advocacy campaigns should frame the regulation and ban of LAWS in terms of well-known cultural symbols that illustrate the familiar fear of technology running amok. One such example is the movie *The Terminator*, where a cyborg killer, the Terminator, is sent back in time to kill all women named “Sarah Connor” in a particular area (*The Terminator* 1984). The website of a current campaign against LAWS called “Campaign to Stop Killer Robots” (Figure 1 and Figure 2) displays only dry and technical information about LAWS that has little appeal to those unfamiliar with the subject. Figure 3 uses the Terminator as a reference point to illustrate the dangers of LAWS. The last point of the section “What Similarities Do Killer Robots Share with the Terminator?” alludes to an iconic scene in the movie (*Ibid.*). The “Do you know?” section informs the reader about current United States internal policy on LAWS and its inadequacy (Department of Defense 2012, 3; Gubrud 2015). Using cultural symbols such as the Terminator, activists can frame the issue of LAWS not as a technical issue of weapon regulation or moral philosophy, but as an issue of unaccountable and unconstrained use of potentially dangerous technology.

Reframing the same issue can be effective in several ways. First, there is already a well-established norm against the use of Terminator-like technology. A senior Department of Defense official, for example, emphasized that the autonomous weapons under development are less like the Terminator and more like the Iron Man from the eponymous movie (Rosenberg and Markoff 2016). There is no need to emphasize the dissimilarity between LAWS and the Terminator unless a norm against the development of technology similar to the latter already exists. Thus, this paper’s proposed frame follows the “grafting” strategy in issue framing (Carpenter 2007, 104). Second, the proposed frame casts the issue in a morally unambiguous way. It also allows the activists to emphasize and illustrate aspects of LAWS that public opinion polls have shown to be the most unnerving, such as circumvention of human decision-making and the machines’ lack of moral conscience (Carpenter 2014). Third, as Figure 3 illustrates, the use of cultural symbols can make technical information memorable and easy to understand. Fourth, just as Skynet is “responsible” for the creation of the Terminator, the proposed frame identifies the governments as the responsible parties for the issue (Keck and Sikkink 1998, 19). Finally, just as the movie is about the potential danger of misused technology, the proposed frame emphasizes the need for regulation and invites concerned citizens to take action against their states’ continued development of LAWS and unwillingness to commit to a regulatory framework.⁴

One objection is that such framing is similar to a scare tactic. Intuition may differ here as to whether this is the case. If the goal is to inform the public in a creative and memorable way, the use of popular culture references should not be problematic. It is important to note that

⁴ Much of the discussion here is U.S. based, mainly because of the importance of the United States in international decision-making, the advanced level of technology development in the United States, and the fact that the United States is the only country so far to have a written guideline on LAWS. See Vilmer, “Autonomous Weapon Diplomacy.” The same strategy, however, applies to all western democracies where there is a cultural norm against similar technology.

calling LAWS “killer robots” is itself an instance of framing, and public opinion research on LAWS shows that the label “killer robots” itself does not make the respondents more opposed to LAWS. It is the idea of LAWS that is feared (Carpenter 2013). Framing is primarily a tool of persuasion and illustration rather than manipulation. While survey research shows that respondents often react unfavorably toward the use of LAWS after the concept is explained clearly to them (Ibid.; Open Roboethics Initiative 2015, 5), this framing technique aims to increase awareness about this highly technical issue and galvanize the public into action.

One may also argue that the proposed strategy is better suited for a campaign that aims for a complete ban on LAWS rather than a seemingly more compromised position of “ban + regulation,” since the proposed frame seems to appeal primarily to the public’s fear of certain technologies. Even if it is the case that the strategy will cause people to push for a complete ban, international activists can capitalize on that momentum to advocate stricter sanctions and regulations on the use of technology than are otherwise possible, especially when partial bans on certain clearly dangerous weapons are met with political resistance from nations. As such, the end result would ideally look very similar to what the “ban + regulation” model prescribes.

On the international level, activists should continue to persuade and offer expertise on the subject. Despite the observation that international norms usually begin as domestic norms (Finnemore and Sikkink 1998, 893), activists need not only focus on the domestic level, especially since CCW meetings are open to experts from NGOs, and there is already a common understanding among most participants of the third CCW meeting that LAWS can be potentially dangerous (Vilmer 2016). The inability to achieve a binding framework or a preventative ban reflects not only the lack of political will, but also the genuine technical difficulty of the subject. Activists and experts of LAWS can provide information and propose a regulatory framework to facilitate the discussion. The following are some suggestions. First, activists should remind participant countries that it may still be too early to settle on any particular definition of LAWS, and activists should try to push the discussion toward a more constructive direction. Second, activists should continue to advocate that LAWS should be subject to “meaningful human control,” where human control takes place or is possible before, during, and after LAWS is deployed (Roff and Moyes 2016). Activists should emphasize that only humans can and should be responsible for the actions of LAWS. Third, activists should highlight certain similarities between LAWS and landmines and “frame” certain regulations of LAWS in terms of landmine regulations. This can lead to a helpful shift from the unfruitful discussion of whether all LAWS should be banned to the identification of the basic elements of a regulatory element. Fourth, activists should advocate for some method of inspection and investigation similar to nuclear inspection, so that AI experts, human rights and extrajudicial killings experts, and military experts can verify that states are not abusing AI technology or violating international law in the development of LAWS. Finally, as more information becomes available and the unacceptability of particular types of LAWS becomes well understood, activists should advocate for the ban of these types of weapons.

Bibliography

- Arms Control Association. "U.S.-Russia Nuclear Arms Control Agreements at a Glance." *Arms Control Association*. April 1, 2014. <https://www.armscontrol.org/print/2556>.
- Asaro, Peter. "On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making." *International Review of the Red Cross* 94, no. 886 (2012). 687-709. doi: 10.1017/S1816383112000768.
- Bryden, Alan. *International Law, Politics, and Inhuman Weapons: The Effectiveness of Global Landmine Regimes*. Abingdon, Oxon: Routledge, 2013. Quoted in John Lewis, "The Case for Regulating Fully Autonomous Weapons," *Yale Law Journal* 124, no. 4 (2015), 1318, <http://www.yalelawjournal.org/comment/the-case-for-regulating-fully-autonomous-weapons>.
- Campaign to Stop Killer Robots. "Learn." *Campaign to Stop Killer Robots*. 2015. <https://www.stopkillerrobots.org/learn/>.
- Campaign to Stop Killer Robots. "The Problem." *Campaign to Stop Killer Robots*. 2015. <https://www.stopkillerrobots.org/the-problem/>.
- Carpenter, Charli R. "Setting the Advocacy Agenda: Theorizing Issue Emergence and Nonemergence in Transnational Advocacy Networks." *International Studies Quarterly* 51, no. 1 (2007). 99-120. doi: 10.1111/j.1468-2478.2007.00441.x.
- Carpenter, Charli R. "How scared are people of 'killer robots' and why does it matter?" *openDemocracy*. July 4, 2013. <https://www.opendemocracy.net/charli-carpenter/how-scared-are-people-of-%E2%80%9Ckiller-robots%E2%80%9D-and-why-does-it-matter>.
- Carpenter, Charli R. "Who's Afraid of Killer Robots? (and why)." *Washington Post*, May 30, 2014. <https://www.washingtonpost.com>.
- Dennett, Daniel C. "When HAL Kills, Who's to Blame? Computer Ethics." In *HAL's Legacy: 2001's Computer as Dream and Reality*. Edited by David G. Stork. 351-365. Cambridge, MA: MIT Press, 1996.

- Department of Defense. *Department of Defense Directive 3000.09: Autonomy in Weapon Systems*. Arlington, VA: United States Department of Defense, 2012. <https://www.hsdl.org/?abstract&did=726163>.
- Finnemore, Martha, and Kathryn Sikkink. "International Norm Dynamics and Political Change." *International Organization* 53, no. 4 (1998). 887-917. <https://www.jstor.org/stable/2601361>.
- "From Not Working to Neural Networking." *The Economist*, June 25, 2016. <http://www.economist.com>.
- Future of Life Institute. "Autonomous Weapons: An Open Letter from AI & Robotics Researchers." *Future of Life Institute*, July 28, 2015. Accessed December 7, 2016. <https://futureoflife.org/open-letter-autonomous-weapons/>.
- Gelly, Sylvain, Levente Kocsis, Marc Schoenauer, Michèle Sebag, David Silver, Csaba Szepesvári, and Olivier Teytaud. "The Grand Challenge of Computer Go: Monte Carlo Tree Search and Extensions." *Communications of the ACM* 55 no. 3. 106-113. doi: 10.1145/2093548.2093574.
- Gibney, Elizabeth. "Google AI Algorithm Masters Ancient Game of Go: Deep-Learning Software Defeats Human Professional for the First Time." *Nature* 529, no. 7587 (2016). 445-446. doi: 10.1038/529445a.
- Gubrud, Mark. "Semi-Autonomous and on Their Own: Killer Robots in Plato's Cave." *Bulletin of the Atomic Scientists*. April 12, 2015. <http://thebulletin.org>.
- Human Rights Watch. *Mind the Gap: The Lack of Accountability for Killer Robots*. New York: Human Rights Watch, 2015.
- Keck, Margaret and Kathryn Sikkink. *Activists Beyond Borders*. Ithaca, NY: Cornell University Press, 1998.
- Lewis, John. "The Case for Regulating Fully Autonomous Weapons." *Yale Law Journal* 124, no. 4 (2015). 1309-1325. <http://www.yalelawjournal.org/comment/the-case-for-regulating-fully-autonomous-weapons>.
- Marshall, John. "Blinding Laser Weapons: Still Available on the Battlefield." *The BMJ* 315 (1997). 1392. doi: 10.1136/bmj.315.7120.1392.

McCormick, Ty. "Lethal Autonomy: A Short History." *Foreign Policy*, January 24, 2014. <https://www.foreignpolicy.com>.

Open Roboethics Initiative. *The Ethics and Governance of Lethal Autonomous Weapons Systems: An International Public Opinion Poll*. Vancouver: Open Roboethics Initiative, 2015.

Owen, Taylor. Lecture through Skype call. Swarthmore, PA: Swarthmore College. November 20, 2016.

"Protocol on Prohibitions or Restrictions on the Use of Mines, Booby-Traps and Other Devices as Amended on 3 May 1996 (Protocol II, as Amended on 3 May 1996) Annexed to the Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May Be Deemed to Be Excessively Injurious or to Have Indiscriminate Effects." May 3, 1996. *United Nations Treaty Series: Treaties and International Agreements Registered or Filed or Recorded with the Secretariat of the United Nations* 2048 (2001): 93-225. <https://treaties.un.org/doc/Publication/UNTS/Volume%202048/v2048.pdf>.

Roff, Heather. Quoted in Ariel Conn, "The Problem of Defining Autonomous Weapons," *Future of Life Institute*, November 30, 2016, <https://futureoflife.org/2016/11/30/problem-defining-autonomous-weapons/>.

Roff, Heather. "The Strategic Robot Problem: Lethal Autonomous Weapons in War." *Journal of Military Ethics* 13, no. 3 (2014). 211-227. doi: 10.1080/15027570.2014.975010.

Roff, Heather. "To Ban or to Regulate Autonomous Weapons: A US Response." *Bulletin of the Atomic Scientists* 72, no. 2 (2016). 122-124. doi: 10.1080/00963402.2016.1145920.

Roff, Heather, and Richard Moyes. "Meaningful Human Control, Artificial Intelligence and Autonomous Weapons." Briefing paper prepared for the Informal Meeting of Experts on Lethal Autonomous Weapons Systems, UN Convention on Conventional Weapons, April 2016. <http://www.article36.org/wp-content/uploads/2016/04/MHC-AI-and-AWS-FINAL.pdf>.

Rosenberg, Matthew and John Markoff. "The Pentagon's 'Terminator Conundrum': Robots That Could Kill on Their Own." *New York Times*, October 25, 2016. <https://www.nytimes.com>.

- Sagan, Scott D. "Why Do States Build Nuclear Weapons?" *International Security* 21, no. 3 (1996/97). 54-86. <https://www.jstor.org/stable/2539273>.
- Sayler, Kelley. "More of the Same: The UN Debate on Lethal Robots." *Center for the Study of the Drone at Bard College*. April 27, 2016. <http://dronecenter.bard.edu/more-of-the-same-the-un-debate-on-lethal-robots/>.
- Tannewald, Nina. "The Nuclear Taboo: The United States and the Normative Basis of Nuclear Non-Use." *International Organization* 53, no. 3. 433-468. <http://www.jstor.org/stable/2601286>.
- Tarrow, Sidney. *The New Transnational Activism*. New York: Cambridge University Press, 2005.
- The Terminator*. Directed by James Cameron. 1984. Beverly Hills, CA: MGM, 2014. DVD.
- "Treaty on the Non-Proliferation of Nuclear Weapons." July 1, 1968. *United Nations Treaty Series: Treaties and International Agreements Registered or Filed or Recorded with the Secretariat of the United Nations* 729 (1974): 161-299. <https://treaties.un.org/doc/Publication/UNTS/Volume%20729/v729.pdf>.
- United Nations Human Rights Council. Twenty-Third Session, Agenda Item 3. *Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions*. Prepared by Christof Heyns. A/HRC/23/47, 2013. http://www.ohchr.org/Documents/HRBodies/HRCouncil/RegularSession/Session23/A-HRC-23-47_en.pdf.
- Vilmer, Jean-Baptiste Jeangène. "Autonomous Weapon Diplomacy: The Geneva Debates." *Ethics & International Affairs*. Carnegie Council for Ethics in International Affairs. September 27, 2016. <https://www.ethicsandinternationalaffairs.org/2016/autonomous-weapon-diplomacy-geneva-debates/>.
- Waltz, Kenneth. *Theory of International Politics*. McGraw-Hill, 1979.
- Wendt, Alexander. "Anarchy is What States Make of It: The Social Construction of Power Politics." *International Organization* 46, no. 2 (1992). 391-425. <https://www.jstor.org/stable/2706858>.