

Killer Robots and Human Dignity

Daniel Lim

Duke Kunshan University
daniel.lim672@dukekunshan.edu.cn

Abstract

Lethal Autonomous Weapon Systems (LAWS) have become the center of an internationally relevant ethical debate. Deontological arguments based on putative legal compliance failures and the creation of accountability gaps along with wide consequentialist arguments based on factors like the ease of engaging in wars have been leveraged by a number of different states and organizations to try and reach global consensus on a ban of LAWS. This paper will focus on one strand of deontological arguments—ones based on human dignity. Merely asserting that LAWS pose a threat to human dignity would be question begging. Independent evidence based on a morally relevant distinction between humans and LAWS is needed. There are at least four reasons to think that the capacity for emotion cannot be a morally relevant distinction. First, if the concept of human dignity is given a subjective definition, whether or not lethal force is administered by humans or LAWS seems to be irrelevant. Second, it is far from clear that human combatants either have the relevant capacity for emotion or that the capacity is exercised in the relevant circumstances. Third, the capacity for emotion can actually be an impediment to the exercising of a combatant's ability to treat an enemy respectfully. Fourth, there is strong inductive evidence to believe that any capacity, when sufficiently well described, can be carried out by artificially intelligent programs.

Background

The Convention for Certain Conventional Weapons (CCW), organized under the United Nations in the 1980s, has sought to restrict or prohibit the use of weapons that are deemed excessively injurious or whose effects are deemed indiscriminate. The convention has successfully levied international restrictions on weapons ranging from landmines to incendiary weapons. More recently, the CCW has added lethal autonomous weapon systems (LAWS), or killer robots, to their agenda. Though there is a lively debate over definitions I will use Wikipedia (https://en.wikipedia.org/wiki/Lethal_autonomous_weapon) to put a working definition of LAWS on the table:

Lethal autonomous weapons (LAWS) are a type of autonomous military robot that can independently search and engage targets based on programmed constraints and descriptions. LAWS are also called lethal autonomous weapon systems (LAWS), lethal autonomous robots (LAR), robotic weapons, or killer robots.

The key phrase in this definition is: 'independently search and engage targets.' LAWS are independent in that they can operate without direct and immediate human oversight. Moreover, they not only *search* for targets they also *engage* targets. That is, they deploy lethal force.

In 2014, signifying its growing prominence on the international stage, the CCW established a Group of Governmental Experts (GGE), involving Russia, China, and the U.S. among others, to investigate the ethical dimensions of emerging technologies relevant to the development of LAWS. ([https://www.unog.ch/80256EE600585943/\(http-Pages\)/7C335E71DFCB29D1C1258243003E8724?Open-Document](https://www.unog.ch/80256EE600585943/(http-Pages)/7C335E71DFCB29D1C1258243003E8724?Open-Document))

A diverse group of voices has been raised against the development of LAWS. In 2015, more than 1,000 AI researchers signed and published an open letter from the Future of Life Institute calling for a ban on LAWS. (<https://futureof-life.org/open-letter-autonomous-weapons/>) Included in this list are well-known technologists like Elon Musk. In 2018, at two of the most recent meetings of the GGE, a number of non-governmental actors made statements in support of banning LAWS. Here is an excerpt from a statement given by the Holy See:

In this regard, the Holy See wishes to remind that a classic foundation of legal systems is the recognition of the human person as a responsible subject that could be sanctioned for his/her wrongdoing and be obliged to provide redress for the damage caused. This notion of responsibility originates from the profound reality of the human person as a free and rational being. An autonomous system has no intention. It merely implements and elaborates algorithms but it does not possess

intentions as such. Autonomous weapons systems that hide the accountable and responsible subject are unacceptable. (Holy See, 2018)

Academic think tanks, like the Center for the Study of Existential Risk at Cambridge University (<https://www.cser.ac.uk/>), have also echoed similar sentiments against the development of LAWS.

Some Arguments

A non-negligible portion of the public, academia, and industry are united in resisting the development of LAWS. They have leveraged a variety of arguments in support of their position. These arguments might usefully be divided along deontological and consequentialist lines.

Many deontological arguments are based on legal compliance failures. Some, relying on International Humanitarian Law (IHL), claim that LAWS do not meet the requirements for distinction (distinguishing civilians from combatants) and proportionality (the force of attacks must properly balance the military advantage it will produce against the civilian damage it will incur). Daniele Amoroso and Guglielmo Tamburrini argue that:

[LAWS must be] capable of respecting the principles of distinction and proportionality at least as well as a competent and conscientious human soldier. (2016, p. 6)

At present, of course, LAWS are not as good as human soldiers in most (if not all) situations.

Of course, many will agree that LAWS, based on the current state of technology, fail to meet these requirements. But this seems to be beside the point since nobody is arguing that LAWS that fail to meet these requirements should be deployed. Consequently, this can only serve as a contingent reason for banning LAWS. Besides, betting against technology, is more often than not, a losing strategy. Recent history has shown, time and again, that machines can encroach on and eventually surpass human expertise when domains are clearly specified. Human world champions in Chess, Jeopardy, and Go can all attest to this. The interesting question is what we should do about LAWS that *do* meet the requirements of distinction and proportionality.

Other deontological arguments have focused on International Criminal Law and the so-called accountability (or responsibility) gap that LAWS create because decision making regarding the use of force will be carried out by non-human entities.¹ Since we don't, at present, have clarity over the concept of non-human accountability, war crimes may

lack legal accountability. Who exactly are we supposed to hold accountable for fatalities caused by LAWS? The absence of accountability, some argue, will potentially be a violation of International Criminal Law.

Thomas Simpson and Vincent Müller's (2015) discussion of accountability is helpful in this regard. They point out that non-military accountability gaps have already been created by other technologies and are deemed acceptable so long as tolerable thresholds of risk are met. It's difficult to see why analogous reasoning might not usefully be applied to LAWS. Consequently, deontological arguments of this sort may also be contingent on technological developments that can help LAWS cross internationally recognized risk thresholds.

Consequentialist arguments for banning LAWS are not explicitly concerned with legal compliance, they are focused on the balance of good and bad outcomes. 'Narrow' consequentialist arguments focus on battlefield performance outcomes: reduced casualties (military and civilian), increased accuracy in targeting, and freedom from self-preservation. Some have also pointed to massive reductions in financial costs. 'Wide' consequentialist arguments focus on geopolitical outcomes: lowered resistance to starting wars, erosion of democratic processes, proliferation of LAWS to oppressive regimes or terrorist groups, and destabilization on a global scale. Narrow consequentialist arguments are often used in support of developing LAWS while wide consequentialist arguments are often used in support of banning LAWS. A full discussion of the relevant consequentialist factors, however, is beyond the scope of this paper.

So much for this all-too-brief look at some of the arguments that have featured in debates over LAWS. This paper will focus on deontological arguments for banning LAWS based on human dignity. Dignity-based arguments are notable because they are not contingent on the state of technology. The absence of contingency is captured well by Peter Asaro:

Any automated process, *however good it might be*, and even if measurably better than human performance, ought to be subject to human review before it can legitimately initiate the use of lethal force. This is clearly technologically required for the foreseeable future because autonomous systems will not reach human levels of performance for some time to come. But more importantly, this is a moral requirement... there is a duty not to permit autonomous systems to initiate lethal force without direct human supervision and control. (Asaro 2012, p. 702, my emphasis)

¹ An anonymous reviewer noted that the class of non-human actors can also include animals and that this fact should be taken into consideration. While this indeed complicates the issue at hand, since international criminal tri-

bunals have limited their jurisdiction to human beings, *any* non-human actors can potentially create problems. Hence, my *not* dealing with animals does not affect my overall argument since all that matters, in this respect, is that LAWS are non-human.

As Asaro points out, it doesn't matter how good LAWS get. Even if they perform measurably better than humans they must not initiate lethal force without direct human control. For Asaro, it is a moral requirement that only a human can initiate lethal force. In other words, truly autonomous lethal weapons should never be developed or used regardless of its performance.² It is my contention, however, that these arguments fail.

The Argument from Human Dignity

Let me begin with a sampling of dignity-based arguments for restricting or prohibiting the development of LAWS. In their general statement at the 2014 Convention on Certain Conventional Weapons expert meeting on LAWS, Germany argued that:

Human control is the foundation of the entire international humanitarian law. It is based on the right to life, on the one hand, and on the right to *dignity*, on the other. Even in times of war, human beings cannot be made simple objects of machine action. (Germany 2014)

The idea here is that allowing machines to make lethal autonomous decisions regarding human combatants is to turn these combatants into the 'simple objects of machine action.' This would rob these combatants of their dignity. This sentiment is also found among academics:

Delegating the decision to kill a human to an algorithm in a machine, which is not responsible for its actions in any meaningful ethical sense, can arguably be understood to be an infringement on *basic human dignity*, representing what in moral philosophy is known as a *malum in se*, a wrong in itself. (Sauer 2016, p. 10)

Call the following argument for a ban on LAWS the Argument from Dignity.

1. If X is killed by a LAWS then X is killed disrespectfully.³
2. If X is killed disrespectfully then X's killing is wrong.
3. If X is killed by a LAWS then X's killing is wrong.

Here I use the word 'disrespectful' as shorthand to describe killings that violate human dignity. This argument is clearly valid, but the truth of premise 1 is questionable. As stated, it unhelpfully begs the question. Michael Horowitz nicely points this out:

While in an esoteric sense, the idea that there is something undignified about dying at the hands of a machine

resonates, why is being shot through the head or heart and instantly killed by a machine necessarily worse than being bludgeoned, lit on fire, or killed by a cruise missile strike? (Horowitz 2016, p. 33)

What we need is an independent reason to believe that the antecedent really entails the consequent in premise 1. One might, following Asaro (2012), argue that:

The very nature of IHL [International Humanitarian Law], which was designed to govern the conduct of *humans* and *human* organizations in armed conflict, presupposes that combatants will be human agents. (Asaro 2012, p. 700, my emphasis)

This sentiment may be regimented through the following two conditionals:

- 1A*. If X is killed by a LAWS then X is killed by a non-human.
1B*. If X is killed by a non-human then X is killed disrespectfully.

Premises 1A* and 1B* can be used to replace premise 1. The basic idea, then, is that the international laws that govern war, like the IHL, can only be applied to human combatants. On its own, however, this claim comes off as *ad hoc* and is no better than baldly asserting premise 1. As Asaro freely admits, this is a decidedly anthropocentric way of making his case. It's not enough, therefore, to simply restrict morally permissible combat to humans. We want to know what it is about being human that makes combat between humans morally permissible in a way that combat between humans and non-humans is *not* morally permissible. In short, we need to know what capacity humans have (that non-humans, and more specifically LAWS, fail to have) that make killings by non-humans disrespectful.

Emotion

A natural place to turn for the critical capacity that makes killing respectful is emotion. Aaron Johnson and Sidney Axinn (2013) develop this idea as follows:

Moral commands are based on values, and values are produced and indicated by sacrifices. While robots are aware of following orders, they are not aware of making sacrifices. Artificial Intelligence still has no real notion of sacrifice. Therefore robots have no values of their own, although they are following the values of their programmers. (Johnson and Axinn 2013, p. 135)

Similar sentiments can be found in Christof Heyns (2016):

truth of 1A* was threatened by the following conditional: if X is killed by a non-human, then X is killed by a LAWS or a non-LAWS (let's call this conditional C). Others might be tempted to draw a similar conclusion. But it seems pretty clear to me that this would be a mistake since the truth of C does not, in any way, threaten the truth of 1A*. Besides, 1A* is arguably a definitional tautology.

² However, it is arguable that there are a number of artifacts already in existence that should be classified as LAWS. These include the SGR-A1 sentry gun and the IAI Harop loitering munition.

³ An anonymous reviewer has suggested that premises 1, 1A*, and 1A^ (among others) can be strengthened by taking other non-human entities (e.g. animals) into consideration. It was even claimed, for example, that the

A machine, bloodless and without morality or mortality, cannot fathom the significance of using force against a human being and cannot do justice to the gravity of the decision. Each instance where force is used against a human being requires that another human being should decide afresh whether to cross that threshold. (Heyns 2016, p. 370)

But, what does it mean to have a ‘real notion of sacrifice’ or to ‘fathom the significance of using force against a human’? The answer to these questions arguably flows from a subsequent point that Johnson and Axinn (2013) make:

[Robots] are not enraged, as humans may be, by the killing of their buddies. But having no emotions, they do not have the attitude toward people that ‘healthy’ humans are expected to have. They do not realize the enormity of an error in killing the ‘wrong’ person. (Johnson and Axinn 2013, p. 136)

LAWS do not have a capacity to have a notion of sacrifice or a capacity to fathom the significance of using force against a human because they don’t feel anything—they have no emotions. It is only because humans can feel the rage and agony that accompanies the killing of humans that they can understand sacrifice and the use of force against a human. Only then can they realize the ‘gravity of the decision’ to kill. So maybe the following two conditionals capture the relevant capacity and can be used to replace premises 1A* and 1B*:

- 1A[^]. If X is killed by a LAWS then X is killed by a combatant that has no emotion.
- 1B[^]. If X is killed by a combatant that has no emotion then X is killed disrespectfully.

Without feeling the pain of losing a friend or feeling the risk of putting one’s life on the line it is impossible to understand the gravity of taking another person’s life.

Let me now state the Argument from Dignity in its developed form based on the human capacity for emotion:

- 1A[^]. If X is killed by a LAWS then X is killed by a combatant that has no emotion.
- 1B[^]. If X is killed by a combatant that has no emotion then X is killed disrespectfully.
- 2. If X is killed disrespectfully then X’s killing is wrong.
- 3. If X is killed by a LAWS then X’s killing is wrong.

4 Responses to the Argument from Dignity

So much for the Argument from Dignity. There are a number of things that might be said in response to it. First, one might follow Dieter Birnbacher and argue that capacities specific to humans (and missing in LAWS) are irrelevant to the respectful killing of humans:

This is spurious. Of course, machines cannot comprehend the value of human life. But why should this make a difference to their victims if alternatively they are threatened to be wounded or killed by manned weapons like bombers? For the victims whose dignity is at stake it is a matter of indifference. (Birnbacher 2016)

Because Birnbacher’s notion of human dignity is based on a given person’s *subjective* experience of humiliation, it is natural for him to think that whether the killer comprehends the value of human life or not is beside the point. What matters is whether the victim experiences a sense of humiliation in the process of getting killed. Victims being threatened with a potential bombing will not care whether the bomb is dropped by a human or a robot.

While this is an interesting response, I’m concerned that there are objective forms of humiliation. This was poignantly depicted in *The Truman Show*, a film about a person named Truman whose entire life was meticulously constructed in order to serve as the centerpiece of a television show for others to watch. While Truman had no negative subjective experiences, there was an objective sense in which he was being used, even humiliated by unknowingly participating in a manufactured existence. Whether the features of Truman’s life are natural or manufactured matter to his dignity. It seems to me that the features of a killer that pertain to the having of certain capacities, like the features of Truman’s life, may matter to a victim’s dignity.

Second, it is far from clear that human combatants meet the necessary emotion-based requirements in order to kill with respect. So a ban on LAWS could *ipso facto* constitute a ban on human combatants as well. Not all human combatants have the capacity to form the relevant notion of sacrifice or the relevant emotions that accompany risk. Humans lie on a spectrum of conceptual and emotional competence. Does this suggest that some killings are more respectful than others?

Moreover, the mere possession of conceptual and emotional capacities cannot be enough. What is needed is the appropriate *exercise* of these capacities. But can human combatants be expected to routinely meet this requirement? This is an extremely demanding request. Human combatants, especially in the heat of battle, hardly have the time or mental / emotional space to exercise the concept of sacrifice or generate the relevant emotions to make informed decisions each time they deploy lethal force. The battlefield, in my opinion, is arguably the worst possible environment to expect something along these lines.

Third, feeling and emotion have a dark side as well. They have the propensity to yield decidedly negative results. Lisa Feldman Barrett, a psychologist who does research on the formation of emotions, writes:

Affective realism can also lead to tragic consequences. In July 2007, an American gunner aboard an Apache

helicopter in Iraq mistakenly killed a group of eleven unarmed people, including several Reuters photojournalists. The soldier had misjudged a journalist's camera to be a gun. One explanation for this incident is that affective realism caused the soldier, in the heat of the moment, to imbue a neutral object with unpleasant valence. Every day, soldiers must make quick decisions about other people, whether they are embedded in a unit during wartime, on a peacekeeping mission, negotiating in a cross-cultural setting... these quick judgments are extremely difficult to negotiate, especially in such high-stakes, high-arousal settings where errors are often made at the expense of someone's life... A little closer to home, affective realism may also play a role in police shootings of unarmed civilians. The U.S. Department of Justice analyzed shootings by Philadelphia police officers between 2007 and 2013 and found that 15 percent of the victims were unarmed... the human brain is wired for this sort of delusion, in part because moment-to-moment interoception infuses us with affect, which we then use as evidence about the world. (Barrett 2017, p. 76)

Being forced to make 'quick decisions' in the 'heat of the moment' soldiers can imbue a neutral object with unpleasant valence and their emotions can generate indiscriminate disrespectful killings. Instead of killing as the result of an emotional process that recognizes the gravity of the situation, human combatants may be emotionally overwhelmed by an instinct to survive. Surely, killings in service of self-preservation can hardly be considered respectful of the victim. The victim, after all, is seen as little more than an obstacle to one's survival.

It is also instructive to take note of Barrett's final quoted sentence. This sort of emotionally charged activity is not, according to Barrett, merely an aberration of human cognition. It turns out the human brain is wired for the kind of delusions that can lead to disrespectful killings (military and civilian) because human action is continuously infused with affect. A case, therefore, can be made (and has been made) that the more emotions are removed from the battlefield, the safer the battlefield can become.

Fourth, there is an unwarranted biological chauvinism regarding emotion that is implicit in this way of reasoning. There are no convincing theoretical reasons why machines cannot have emotions. At the very least, because we have to acknowledge the existence of a lively debate on this topic, we cannot assume that emotions are unique to human biological systems. In fact, there seems to be a growing majority of philosophers and scientists that believe emotion (or consciousness in general) is multiply realizable and substrate independent.

Hilary Putnam (1967), a philosopher, noted that a wide variety of terrestrial creatures are capable of experiencing pain. If the mental state of experiencing pain were identified with a specific physical-chemical state of the *human* brain

then this state must also be a possible state of a reptilian brain or a mollusk's brain if (all should agree) reptiles and mollusks are to be able to experience pain. At the same time, it must not be a possible state of any physically possible creature that cannot experience pain. The requirement that mental states be limited to human biology is far too restrictive, and in Putnam's view, must be relaxed. He rightly defends the claim that mental states must be realizable in multiple physical ways. There is a certain level of independence that mental states have with respect to the physical systems they are realized in.

Max Tegmark, a physicist, echoes these same sentiments:

In other words, the hardware is the matter and the software is the pattern. This substrate independence of computation implies that AI is possible: intelligence doesn't require flesh, blood or carbon atoms. (Tegmark 2017, p. 67)

While Tegmark, in this quote, is specifically referring to the possibility of intelligence in non-biological systems, he goes on to write in support of an intelligence-based theory of consciousness (and hence emotion) that relies on integrated information (a doubly independent abstraction above physical matter).

Finally, Barrett makes essentially the same observation:

A mental event, such as fear, is not created by only one set of neurons. Instead, combinations of different neurons can create instances of fear. Neuroscientists call this principle degeneracy. Degeneracy means 'many to one': many combinations of neurons can produce the same outcome. In the quest to map emotion fingerprints in the brain, degeneracy is a humbling reality check. (Barrett 2017, p. 19)

So there seem to be no principled reasons to think that mental states, like the ones being called on for making respectful killings possible, cannot be realized in LAWS. And if, following Barrett (2017), we conceive of emotions as analogues of perceptions then, given our expected technical mastery over perception, there seem to be no essential roadblocks to our technical mastery over emotions. If emotions are really nothing more than simulations of one's *internal* body in the way that perceptions are simulations of one's *external* world, then, the project of imbuing AI with emotions fits squarely within the realm of possibility.

Tegmark's brief allusion to theories of flight may offer additional help. If we limited our concept of flight to what was available to our ancestors, the *avian* variety, we would've been forced to treat all of our more recent breakthroughs in aviation technologies as something other than genuine flight. We would have been tempted to say that our airplanes don't 'really' fly, they merely simulate flight. But we know better than to give in to this kind of avian-centric thought. Flight is not limited to the flapping of wings. The

flapping of wings is just one way of realizing flight. Flight, it turns out, is multiply realizable.

Similarly, emotions fall into the same category. Should we be tempted to say, because the silicon-based hardware used to run AI programs are different from the carbon-based hardware used to run human minds, that AI can't 'really' have emotions, they merely simulate them? It seems to me that such a temptation is guilty of excessive anthropocentrism. There is no *a priori* reason to expect emotion *simpliciter* to be limited to the category of human emotion.

Concluding Remarks

I spent some time discussing the capacity for emotion. There are, no doubt, other candidate capacities that might be used for justifying premise 1 of the Argument from Dignity. For example, following Sparrow (2016) and Nagel (1972) before him, one might argue that the capacity to engage in interpersonal relationships is necessary for killing respectfully:

The relationship is fundamentally a relationship between agents... immediately, then, we can see why [LAWS] might be thought to be morally problematic, regardless of how reliable they might be at distinguishing between legitimate and illegitimate targets. When [LAWS] decide to launch an attack the relevant interpersonal relationship is missing. Indeed, in some fundamental sense, there is no one who decides whether the target of the attack should live or die. The absence of human intention here appears profoundly disrespectful. (Sparrow 2016, p. 107)

This could be used to ground the following two conditionals that might replace premise 1:

- 1A~. If X is killed by a LAWS then X is killed by a combatant without the relevant interpersonal relationship.
- 1B~. If X is killed by a combatant without the relevant interpersonal relationship then X is killed disrespectfully.

What makes interpersonal relationships possible, according to Sparrow, is the presence of human *intention*. Many of the same things said about human emotions above, however, can easily be adapted as responses to human intentions: (i) it is irrelevant with respect to subjective human dignity, (ii) it is seldom, if ever, exercised by human soldiers in actual combat, (iii) it can have adverse consequences, and (iv) it is far from clear that machines cannot be developed to have genuine intentions.

In this brief discussion, I have tried to argue that a non-contingent ban on LAWS based on human dignity fails. This is not to say that a ban on LAWS cannot be justified. As mentioned above, there are other ways (e.g. accountability gaps, wide consequentialist factors) to argue for a ban. But

given the beneficial possibilities that arise out of the advancement of these technologies, it would be detrimental to put a premature moratorium on the development of LAWS.

References

- Amoroso, D., and Tamburrini, G. 2018. The Ethical and Legal Case Against Autonomy in Weapons Systems. *Global Jurist* 18 (1), DOI: 10.1515/gj-2017-0012.
- Asaro, P. 2012. On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making. *International Review of the Red Cross* 94: 687-703.
- Barrett, L. F. 2017. *How Emotions are Made: The Secret Life of the Brain*. New York: Houghton Mifflin Harcourt.
- Birnbacher, D. 2016. Are Autonomous Weapons Systems a Threat to Human Dignity? In N. Buhta (ed.) *Autonomous Weapons Systems: Law, Ethics, and Policy*, pp. 105-121.
- Germany. 2014. General Statement. *Convention of Certain Conventional Weapons Expert Meeting: Lethal Autonomous Weapon Systems*.
- Heyns, C. 2016. Human Rights and the use of Autonomous Weapons Systems (AWS) During Domestic Law Enforcement. *Human Rights Quarterly* 38: 350-378.
- Holy See. 2018. "Statement by the Permanent Mission of the Holy See to the United Nations." ([https://www.unog.ch/80256EDD006B8954/\(httpAssets\)/E3015FACDD8757B4C12582F800595E37/\\$file/2018_GGE+LAWS+2_6a_Holy+See.pdf](https://www.unog.ch/80256EDD006B8954/(httpAssets)/E3015FACDD8757B4C12582F800595E37/$file/2018_GGE+LAWS+2_6a_Holy+See.pdf))
- Horowitz, M. 2016. The Ethics and Morality of Robotic Warfare: Assessing the Debate Over Autonomous Weapons. *Daedalus* 145 (4): 25-36.
- Johnson, A., and Axinn, S. 2013. The Morality of Autonomous Robots. *Journal of Military Ethics* 12 (2): 129-141.
- Nagel, T. 1972. War and Massacre. *Philosophy and Public Affairs* 1 (2): 123-144.
- Putnam, H. 1967. Psychological Predicates. In W. Capitan and D. Merrill (eds.) *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, p. 37-48.
- Sauer, F. 2016. Stopping 'Killer Robots': Why Now is the Time to Ban Autonomous Weapons Systems. *Arms Control Today*: 8-13.
- Simpson, T., and Müller, V. 2016. Just War and Robots' Killings." *Philosophical Quarterly* 66 (263): 302-322.
- Sparrow, R. 2016. Robots and Respect: Assessing the Case Against Autonomous Weapon Systems. *Ethics and International Affairs* 30 (1): 98-116.
- Tegmark, M. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York: Alfred A. Knopf.