Band 19 / 2024

Security through Unity

Europe’s Challenges after Ukraine Crisis

Paul Ertl (Ed.)

Schriftenreihe der
Landesverteidigungsakademie

UNSER HEER

Schriftenreihe der Landesverteidigungsakademie

Paul Ertl (Ed.)

Security through Unity

Europe's Challenges after Ukraine Crisis

19/2024
Vienna, December 2024

Imprint:

Copyright, Production, Publisher:
Republic of Austria / Federal Ministry of Defence
Rossauer Lände 1
1090 Vienna, Austria

Edited by:
National Defence Academy
Center for Leadership and Defence Awareness
Stiftgasse 2a
1070 Vienna, Austria

December 2024
ISBN 978-3-903548-04-6

Printing:
ReproZ W 25-xxxx
Stiftgasse 2a
1070 Vienna, Austria

Table of contents

Autonomous Weapons for Humanity: Lessons from the Russo-Ukrainian War

Nathan Wood

Abstract

Though autonomous weapon systems (AWS) fulfill a number of critical combat roles, and have done so for decades, ethical and legal worries have been raised in relation to recent advances in artificial intelligence (AI) and the possibility for more open-ended or versatile deployments of these systems. In particular, the potential for AWS to be used in anti-personnel roles is argued by many to present especially difficult issues ranging from concerns over the dignity of those targeted by such systems to practical worries that AWS will be incapable of making the nuanced judgments required for distinguishing combatants from civilians. However, the ongoing Russo-Ukrainian War highlights a potential response to these objections, namely that autonomous weapons may be deployed in order to further humanitarian aims by providing a means for would-be interveners to create humanitarian corridors for fleeing civilians. The argument focuses on one of the most objectionable forms of AWS, namely explicitly anti-personnel systems which aim to kill combatants, and argues that even rudimentary AWS with limited abilities of distinction could be used to bolster humanitarian aims while adding only a minimal risk to civilians. This is because object recognition in computer systems has reached a level where such systems would be able to rather reliably determine whether an individual is armed or not (and these capabilities are continually improving), and based on this criterion alone AWS could be deployed to create safe corridors of retreat for fleeing civilians. Such corridors could be marked to all sides in a conflict, their locations communicated to civilians seeking refuge, and the AWS programmed to lethally engage anyone entering the designated corridor who is armed. As the purpose of humanitarian corridors is to enable civilians to safely exit a combat zone, there is no

good reason for armed units of either side to enter such a space, and so targeting based solely on whether or not one is armed could also be used as a means to show resolve in furthering humanitarian aims while taking a decidedly neutral position in the conflict itself. AWS deployed for such purposes could thus engage armed units with far less discrimination, aiming at anyone who is armed, rather than only "enemy" combatants, and the securing of such corridors through wholly unmanned systems has the added benefit that it risks no lives of would-be interveners, making participation more likely. The general simplicity of the targeting parameters for AWS deployed to such tasks (if armed, engage, otherwise, not) would furthermore ensure that the interveners could only be utilizing their autonomous weapons for purely humanitarian aims rather than as a cover for bolstering the position of one side in a conflict. The arguments further highlight a potential use-case for what is widely considered one of the most objectionable forms of AWS, presenting a counterpoint to critics' claims.

Keywords: *Autonomous Weapon Systems, International Humanitarian Law, Distinction, Humanitarian Intervention*

Introduction

Autonomous weapon systems (AWS) have been in use for decades, performing a number of unobjectionable combat roles around the world. However, as artificial intelligence (AI) is incorporated to an ever-greater degree in these systems, creating risks of opaque decision-making and potential unpredictability, critics of autonomous and AI-enabled weapon systems have become increasingly vocal in their objections to these. One of the most widely criticized types of such platforms, and certainly one of those bringing with it the greatest ethical and legal challenges, are lethal anti-personnel systems. In the literature critical of AWS, these are what are most emblematic of the charged terminology of "killer robots".[1] However, as I have argued elsewhere, it is not the

[1] See, e.g., Asaro (2012); Human Rights Watch (2012); Sparrow (2016).

systems themselves that present ethical or legal challenges, but rather uses (and misuses) of those systems.[2] In particular, though almost any system can be used in an unethical or illegal fashion, nearly all see at least some permissible use cases, and some cases where use of these weapons is arguably superior, on not only moral and legal grounds, but also on broader pragmatic or strategic grounds such as deterrence or strategic stability. In this article, I present and defend one such use case for lethal autonomous anti-personnel systems, namely the establishment of humanitarian corridors (HCs) in active war zones, demonstrating that even these platforms most subject to critique (and rightly so) may find situations where their deployment is morally, legally, and strategically sound. Importantly, the argument does not imply that such systems are in general permissible, as there are many contextual factors that impact on the permissibility of a certain weapon's use in discrete warfighting roles. However, it shows that blanket prohibitions may undercut certain uses which are compliant with both the ethics and laws of war, and are moreover of a purely humanitarian nature. The upshot of the arguments is thus that we should not seek overarching positions on AWS, but instead look to each potential system, its potential uses, and the contexts of such use, developing rules and best practices to guide these while recognizing that nearly all weapons, despite their destructive potential, can be used for good or ill.

The arguments are structured as follows. I begin in Section 2 by discussing the importance of humanitarian corridors for protecting civilians fleeing from war, highlighting what is necessary for an HC to be effective, respected by all sides, and properly neutral. In Section 3, I examine what minimal targeting parameters would need to be instituted in order for an AWS-secured HC to meet the requirements sketched in the previous section. As the arguments focus on AWS deployed for humanitarian purposes by states not party to the conflict, Section 4 addresses the potential objection that deploying *lethal* autonomous *weapon* systems may

[2] Wood (2023b); Wood (unpublished manuscript).

expand the conflict, drawing in neutral parties. Finally, Section 5 concludes by making clear that the arguments are not in defense of anti-personnel systems *per se*, but rather are meant to indicate that over-arching positions or blanket prohibitions are apt to be too broad, as there will exist legitimate use cases for almost every type of system we might develop. Moreover, though systems can (and likely will) be misused, good actors can utilize AWS to bolster human rights and respect for life, and we have a responsibility to be mindful of each and every use case which can further these important goals.

However, before moving onto the arguments, there are two points in need of addressing. First, it is critical that we are clear on what we mean by "autonomous weapon system".[3] Though there exist many competing definitions in both academic and governmental positions,[4] there is a growing consensus which follows the views of the United States Department of Defense (DoD) and the International Committee of the Red Cross (ICRC) in treating as AWS all systems which, "once activated, can select and engage targets without further intervention by a human operator".[5] This view, moreover, is closely followed by many critics of AWS, including the civil society group known as the Campaign to Ban Killer Robots.[6] As the arguments to come resist the conclusions of these critical voices, the broad definition of AWS provided by the DoD and ICRC and echoed by critics thus provides a useful departure point.

Second, it is worth stating clearly that the arguments of this article focus exclusively on (potentially) lethal *anti-personnel* autonomous weapon systems. Anti-personnel systems pose the most incisive moral and legal challenges, and beyond this, anti-material systems would seem

[3] See Wood (2023a) for a defense of the need for definitional clarity.
[4] Williams (2015); Boothby (2016); Altmann and Sauer (2017); Caron (2020); Taddeo and Blanchard (2022); Oimann (2023) all provide useful exploration of varying definitions. See also Wood (2023a).
[5] 5US Department of Defense (2023), pp. 21. See also International Committee of the Red Cross (2014), p. 5 and International Committee of the Red Cross (2021), p. 1.
[6] See, e.g., Asaro (2019); Human Rights Watch (2020), p. 2; Stop Killer Robots (2022), p. 6.

to be transparently acceptable when used to prevent incursion into HCs by belligerent armored vehicles or combat aircraft. More than this, HCs ought not have any combat assets within them, and the ease with which anti-material systems can identify vehicles and incoming ordnance that may be intercepted makes clear their permissibility. It is also worth making clear that the autonomous systems being discussed are ones where there is neither contemporaneous human input to the AWS' decisions, nor is there a possibility for a human to override those decisions. Using alternate terminology, the focus is on "human off-the-loop" systems rather than those with a human "in-the-loop" or "on-the-loop". The purpose of this limitation is to target the discussion to the most morally and legally challenging types of AWS, namely anti-personnel systems where a human cannot contribute to or override the decisions of the system. If the arguments developed below are able to show that even these systems have an ethical and legal use case, then this serves as a particularly strong objection to arguments for partial or total bans of AWS development and deployment.

1 Humanitarian Corridors

Warfare has never been a clean enterprise, but in modern conflict the costs which civilians bear are often truly terrible. More than this, civilians are often intentionally targeted by belligerent parties, and are sometimes used as a moral or legal shield by bad actors who either hide within the civilian population or place civilians between themselves and opposition forces. As cases in point, the ongoing Russian invasion of Ukraine has seen regular and immense bombardment of Ukrainian civilian sites by Russian forces, and the recent Hamas incursion into Israel and subsequent Israeli counterattack highlight the precarious situation of civilians on all sides of conflict. In any war where hostilities take place at least partly within populated areas, or where military forces have clear and legitimate objectives within cities or civilian spaces, there is moral and humanitarian value in doing all that we can to remove civilians from

combat zones before hostilities place them at severe risk.[7] Additionally, international humanitarian law (IHL) requires that certain precautions in attack be taken when conducting hostilities which may be expected to incidentally harm civilians. All of these factors argue for the establishment of humanitarian corridors when combat is expected to spill over into civilian spaces. But how exactly should such corridors be established? More importantly, who is to establish such corridors, through what means, and if fighting is required to protect the civilians fleeing, how is that to be done in a non-escalatory manner which does not risk exacerbating the conflict by drawing more parties into it? The remainder of this section will be devoted to discussing four elements which are required for a humanitarian corridor to be effective and sustainable,[8] and the following section will present a potential means for fulfilling these criteria.

The first and arguably most important element of any effective humanitarian operation, be it the creation of a safe zone, establishment of a corridor through which to flee the fighting, or even the instatement of a temporary ceasefire, is that it is properly protective; merely indicating that an HC exists will not suffice. Nor is it acceptable to have demarcated corridors which are unsafe for civilians or which take civilians into danger. Thus, HCs which lead through mined areas, which are not protected, or which take civilians into hostile zones (or into captivity) are all insufficient. More than this, such corridors would be in breach of both the ethics and laws of war which prohibit subjecting civilians to disproportionate and unnecessary harms. A humanitarian corridor must help civilians to escape the fighting, and it cannot subject them to more risks than would be suffered if they remained where they were.

---

[7] It is likely infeasible for civilians to be evacuated before hostilities commence at all, and it is likewise often tactically problematic to provide warning too far in advance. However, there are many cases where civilians can and should be removed, such as during prolonged sieges of urban areas, or in combat where combatants have dug in to such an extent that significant destruction and use of less accurate munitions is expected during the course of fighting.

[8] This is not to be taken as an exhaustive list of the necessary components, but rathera sketch of the broad elements needed.

Second, in order for an HC to remain effective, it will likely need to be neutral. What this means is that the corridor cannot present any side to a conflict with some clear military advantage. It may not seem immediately clear how this even could be the case, but if one imagines a corridor as an inviolable space, and that corridor is stretched across a state, then this would create an effective barrier to operations, potentially allowing one side to dig in, prepare an assault, or consolidate gains behind the corridor's line of march. As the tempo of war may ebb and flow, such a corridor may be a hindrance to a party one day, a boon the next. But importantly, if the corridor provides any substantive advantages for any significant amount of time, it is likely to be crossed by combat units, risking the chance that fighting erupts within its bounds, and thereby creating grave risks for the civilians traversing it (who may also be more tightly grouped within that space, compounding the risks to them).

Third, civilians within an HC are likely to be at even greater potential danger than if they were in their homes waiting out the fighting. This is because they will often be in the open and presented with much less cover in the event of fighting, they will be traveling in easily distinguishable civilian groups (putting them at risk of strike from criminal belligerents who may seek to actively target them), and they will almost certainly be more densely grouped than if they were each hiding on their own. Because of these factors, it is critical that HCs are robust. What this means is that an HC, once established, has mechanisms for guaranteeing that it remains truly protective for those within it for so long as the corridor is being maintained. The central element of robustness is that the HC is manned or covered by combat units which are prepared to use force to guarantee the safety of civilians. If this is not the case, then the corridor cannot be seen to be truly protective, as a criminal belligerent may strike the civilians there at any moment. Thus, the capacity to use lethal force, and the clear and transparent intention to do so when necessary, are central to building effective HCs in areas where some belligerents refuse to afford civilians their due protections under IHL.

If would-be interveners are prepared to use lethal force to guarantee the safety of those within a humanitarian corridor, a fourth and final element which is critical is that this force (and thus the corridor itself) not prove escalatory or risk drawing additional states or groups into the conflict. It may seem obvious, but if the establishment of an HC presents serious risks of intensifying a conflict and expanding it, then any gains it may provide for discrete civilians fleeing at that moment are apt to be overshadowed by the far larger risks of exacerbating the situation. Thus, humanitarian corridors, while almost certainly requiring guarantees of safety backed by the threat of force, cannot be set up in a manner expected to intensify, extend, or expand the fighting.

2 AWS for Humanity: Design requirements and targeting parameters

Humanitarian corridors must be designed to be truly protective of civilians, robust in the protections they offer, neutral enough to be respected by all sides to conflict,[9] and careful enough to be non-escalatory in nature. One may propose different means of achieving these goals, but in this section we will focus on only one such possibility, namely the use of AWS as defenders of civilians fleeing along an HC. In developing this point, I will focus on the minimal capabilities such AWS must have and the targeting parameters which are necessary in order for the HC to be protective, robust, neutral, and non-escalatory.

In this article, we are focusing on explicitly anti-personnel AWS, but it is worth reiterating that effective HCs will likely require anti-materiel systems as well, either of an autonomous or manned variety, in order to prevent combat vehicles and long-range weapons from being used against civilians. This is especially important if one considers the regu-

[9] This requirement is qualified as HCs are apt to always present at least some passing advantage to one side in a conflict. The critical point is that the advantage not be great enough, or favor one side often enough, that it undermines the general respect for the HC by all sides.

larity of Russian attacks on civilians in Ukraine through the use of artillery and long-range missiles, making clear that any attempt to protect civilians, including those who are fleeing, will require systems to counter these threats. Anti-materiel systems present far fewer ethical or legal hurdles though (targeting mistakes are far less likely in anti-materiel systems, and systems countering artillery or missile batteries which are firing on civilians do not raise special problems). Given this, in discussing what is needed for establishing effective HCs, we will take for granted that anti-materiel systems are in place and focus on what capabilities and targeting parameters must be present in order for anti-personnel systems to also be permissibly deployed.

With regards to design needs and capabilities, anti-personnel AWS must, at a minimum, be able to reliably recognize small arms, be able to be geolocked to certain areas of activity, be capable of targeting and engaging individuals (lethally, if necessary), and be highly accurate in their engagements. The ability to recognize small arms is central to locating individuals who may potentially threaten civilians in the HC, the ability to target and engage such individuals is necessary for the HC to be properly protective, geo-locking allows one to limit the AWS to operations only within the HC (preventing spillover conflict that may be escalatory), and accuracy is required to prevent incidental harms to civilians. Accuracy is indeed doubly important, as not only ought civilians not be harmed by the AWS' operations, but an AWS that collaterally harms some civilians in the protection of others may undermine civilians' overall trust in the AWS and the humanitarian corridor, reducing their reliance on it and limiting its humanitarian value. If an AWS, however, can competently recognize small arms (subject to extensive testing and evaluation to demonstrate robustness in that capacity), can accurately engage armed individuals entering the corridor, and can be suitably limited to only engaging armed individuals within the bounds of the corridor, then it may be given targeting parameters which would allow it to be effectively used, and used in accordance with the four broad requirements of protectiveness, robustness, neutrality, and non-escalation.

Technical capabilities form only a part of responsibly and reliably utilizing AWS for establishing HCs though. In addition to these minimal hard- and software requirements, an AWS must be given targeting parameters which foster and promote the four essential values listed above. One might imagine that these targeting parameters would of necessity be complex, or require a nuanced approach to identifying potential risks to civilians. However, this could not be farther from the truth. If the aim is to create a safe space for civilians, and to do so in a maximally neutral and ideally transparent fashion, then an AWS may be programmed to simply target any and all individuals who are visibly armed and within the bounds of the HC.

These parameters are straightforward, making compliance with them clear for civilians who are entering the corridor. This is critical given that civilians in war zones may often be armed themselves, and it will be necessary that they lay down arms (ideally, at a point where they can be tagged and collected by humanitarian workers or autonomous systems). Moreover, given that geo-location is a standard and reliable capability for modern platforms, the only real task the AWS must carry out is an object recognition exercise in determining whether or not small arms are being carried by any individual. This task is increasingly within the capacity of current systems,[10] and will only become more reliable as further developments are made in AI and AWS. One might object that if the AWS is only recognizing *visible* arms, but not actively searching indviduals entering the corridor, then it would still be possible for bad actors to sneak weapons into the HC and then fire on civilians. Sadly, this is true, and cannot be easily dealt with in a straightforward fashion. However, it bears emphasizing that the moment an individual brandishes arms, they will be targetable for the AWS. Sufficient coverage of the HC by autonomous systems thus might allow for covert infiltration

[10] Winter (2020) provides useful discussion of object recognition capabilities in AWS. See also Scharre and Sayler (2016) and Scharre (2023) for more in-depth discussion of AI capabilities relevant to this task.

by criminally-minded combatants, but it would also guarantee that the moment, they open fire they are likely to be found and engaged.

On a similar note, one might also object that a corridor monitored only for openly carried arms would allow for terrorists or infiltrators to use the HC to sneak into enemy positions or through otherwise secured locations. To this, it is worth stating that while this is a real risk, and not to be discounted or brushed away, it is a risk attending any mass movement of civilians; individuals will always find better opportunity to mask their movements within such groupings, and it will be difficult to root out those utilizing the HC as a source of cover. However, it is also worth stating that it will likely be clear in many cases when individuals *exit* the HC at points other than within refugee camps or transfer facilities, and so many instances of combatants using the corridor as cover may be discovered as they leave it. Moreover, as AWS can rather easily be outfitted with video recording devices, footage of such combatants could be used against them later, as their surreptitious use of civilians for cover would in most cases be in breach of IHL and ground their being tried as war criminals.

Looking only to whether or not individuals are armed provides a simpler targeting profile for AWS, making them likely to be much more reliable in their role as protectors of civilians. However, this extremely bare target profile is also critical for anchoring the value of neutrality in humanitarian corridors. This is because, by targeting anyone who is armed, the HC by default will not be tied to one side in the conflict. Moreover, this matches the neutral humanitarian nature, which such corridors should have, as these should be tracking basic human needs and not bolstering one party to a conflict. A transparent implementation of AWS with such bare target profiles signals to all parties that the HC really is of a purely humanitarian nature, and moreover provides no side with a clear way to make use of or exploit the HC. There will exist larger strategic considerations necessary for ensuring that the HC is properly neutral and non-escalatory, but the basic target profile used by the AWS sets a

foundational element of neutrality critical for long-term success in such humanitarian missions.

3 Strategic Stability and (De-)Escalation

Deploying autonomous weapons with the above target profile may allow interveners to claim, in good faith, that they are acting as neutral parties securing human rights and are not themselves engaging in the conflict. However, it remains true that these AWS are capable of engaging belligerents from all sides of a conflict, and we can expect at least some such engagements to occur. After all, even a party wholly intending to respect the HC will likely have some combatants who mistakenly enter its area and are thereby (lethally) engaged. Such violence risks expansion of the conflict, and for HCs to be truly effective in their humanitarian mission, this objection from escalation must be addressed. More than this, Russia's ongoing invasion of Ukraine has been attended by a series of threats against would-be interveners, either in the form of veiled statements about "consequences... such as you have never seen in your entire history"[11] to regular efforts to inspire fear of nuclear responses against those who might seek to actively support Ukraine's fight of self-defense.[12] These threats have proven to be hollow, and future threats should also be so treated, but even so, a belligerent who treats *any* outside "interference" in such a highly charged manner would be likely to cast efforts to create HCs as somehow direct attacks on their state. As such, it is critical that the establishment of humanitarian corridors through the deployment of AWS be done in such a way as to minimize these expectable risks of escalation and ensure that strategic stability is maintained.

The first point worth strongly reiterating is that properly neutral humanitarian corridors can greatly mitigate the escalatory potential of deploying combat systems to ongoing conflicts. In particular, if the

[11] Excerpt from Putin's February 24th address. Official transcript available at http://en.kremlin.ru/events/president/news/67843.
[12] Pifer (2023).

rather open targeting parameters sketched above are utilized, then *all* sides to conflict can expect their troops to be engaged when armed and entering a corridor. Such impartiality in attack can go a long way to making clear that one is not taking sides in a conflict. More than this, when a party providing the AWS to secure the corridor strongly supports one belligerent actor in the conflict, maintaining this neutrality in targeting is all the more important to ensure that the intervener itself does not violate the humanitarian nature of their mission. This is worth stating, as it would be all too easy for well-meaning rights respecting regimes to subtly alter the programming and target profiles of their AWS to make them less likely to target combatants from the group they support. However, such duplicity in deed would undermine the HC itself, inviting attacks from belligerents who feel it is being used as cover for military intervention. This would place civilians within the corridor at grave risk, and strongly speaks against such actions.

A second factor which is important in the establishment of HCs is transparency, and in order for civilians to make use of the corridor and belligerents to respect it, there are a number of points which must be clearly conveyed to all relevant parties. First, it must be entirely clear what the targeting parameters of the AWS are so that civilians can avoid being targeted. Second, the exact route of the corridor must be given to all belligerents so that they may direct their forces around it. Similarly, civilians must be made aware of the corridor's location so that they know both where they should flee to, and also at what point they must disarm themselves.[13] Third, potential limitations of the AWS deployed must be communicated to belligerents and civilians so that they know what sorts of action or what sorts of items may be apt to lead to targeting mistakes. Fourth, the automaticity of the AWS, and the impossibility of override by humans, must be made unmistakeably clear to belligerents, so that

[13] To the latter point, there would ideally be weapon collection points, where civilians can register their name and weapon with intervener forces so that they could collect their weapons at a later time. Such collection points would moreover provide clear reference points for where the edge of the corridor lays.

they are fully aware that failure to respect the HC will result in guaranteed lethal response. With parties like Russia, this is even more important, as it shifts the escalation ladder back onto them, thereby preventing them from committing breaches in order to "test resolve" of humanitarian actors; by making lethal responses to incursion into the HC a foregone conclusion, those breaching the HC are left with virtually no recourse to claim they were wrongfully attacked, which thereby dampens possibilities of escalation. Furthermore, the clear communication of exactly where the corridor is prevents any potentially escalatory "mistakes" from occurring; combatants know where they should (not) be, and if they enter the HC and are targeted, that is their own fault.

Finally, related to the points on neutrality discussed above, it is likely that HCs will need to be periodically opened to allow combat units to cross them. Scheduled openings with clearly marked "safe zones" where civilians can securely group may serve to ensure that the corridor does not present a persistent barrier to operations, while still allowing it to convey civilians across large swaths of land. Such temporal limits bolster the image of the intervener as a purely humanitarian actor and not an active party in the hostilities. Moreover, it undermines any argument from military necessity a belligerent may give for justifying incursions into the HC outside of scheduled opening times. Ultimately, the regular scheduled opening of the corridor, and in many locations, is necessary for strategic stability as it ensures that the corridor itself does not (greatly) alter the strategic considerations of any belligerent to a conflict.

As a final point, no one should be aiming to strike civilians, and so no one can object on strategic grounds to a platform, which is deployed specifically to secure civilians, and is limited in its operations to the area of a humanitarian corridor. There are clear military grounds for needing to move across the area covered by an HC, especially if that corridor is lengthy or bisects areas of strategic relevance, but the corridor may be opened regularly enough to eliminate even this objection. And while one may argue that an AWS capable of lethally engaging combatants presents serious risks of escalation, there is also an important sense in which this

shows resolve on the part of would-be interveners.[14] In fact, setting an AWS into fully autonomous mode makes clear that any breach of the HC will result in engagement, and will cause this without human agency being brought into the picture at all. Transparent communication of this fact will give potential breachers of the corridor pause, and set them in a poor position to claim they were "wrongfully" targeted by the AWS. These factors conspire to make the AWS potentially more deescalatory than flesh-and-blood combatants might be, and give interveners a means of establishing effective humanitarian corridors without putting their own boots on the ground.

Conclusion

Successfully establishing humanitarian corridors is as much a political task as a military one, and the arguments developed here should be taken as nothing more than a brief sketch of some of the factors involved. However, even from this cursory presentation, it is clear that anti-personnel autonomous weapons may find a legitimate use case in this type of scenario. In particular, anti-personnel AWS present a politically viable, strategically sound, and robust means for would-be interveners to put the needed "boots on the ground" for a corridor to be maintained. Transparent communication of the AWS' exact targeting parameters, the exact location of the corridor, the ways in which one can safely enter and exit it, and the times and places when the corridor will be opened for combat units passing through all show how such corridors may moreover be established to be properly neutral and maximally non-escalatory. Removing civilians from combat zones moreover reduces the overall civilian harms which may be suffered in war, lowering the total grievances a party may have, and increasing the chances of a lasting peace being secured as quickly as possible.

[14] See, e.g., Wong et al. (2020).

In closing, it is worth stating clearly that the presence of a use case for anti-personnel autonomous weapons does not indicate that these would in general be permissible to deploy. The capabilities of a particular system and its particular context of use will greatly impact on its overall permissibility (moral or legal) in discrete scenarios. The critical lesson to be taken away is that we ought not search for overarching arguments for or against (classes of) systems, but instead craft rules and best practices for how systems are used across an array of contexts. This is especially so, given that even the most seemingly morally objectionable types of platform – lethal anti-personnel off-the-loop autonomous weapons – may be used in an ethical and legal manner, and for the purposes of protecting humanity and bolstering the laws of war.

References

Altmann, J. and F. Sauer (2017). Autonomous weapon systems and strategic stability. *Survival 59* (5), 117–142.

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* (886), 687–709.

Asaro, P. (2019). Algorithms of violence: Critical social perspectives on autonomous weapons. *Social Research: An International Quarterly 86* (2), 537– 555.

Boothby, W. H. (2016). *Weapons and the Law of Armed Conflict* (2 ed.). Oxford, UK: Oxford University Press.

Caron, J.-F. (2020). Defining semi-autonomous, automated and autonomous weapon systems in order to understand their ethical challenges. *Digital War 1* (1- 3), 173–177.

Human Rights Watch (2012). Losing humanity: The case against killer robots.

Technical report, Human Rights Watch.

Human Rights Watch (2020). New weapons, proven precedent: Elements of and modes for a treaty on killer robots. Technical report, Human Rights Watch.

International Committee of the Red Cross (2014). Autonomous weapons systems: Technical, military, legal and humanitarian aspects. Technical report, International Committee of the Red Cross.

International Committee of the Red Cross (2021). ICRC position on autonomous weapons systems. Technical report, International Committee of the Red Cross.

Oimann, A.-K. (2023). The responsibility gap and LAWS: A critical mapping of the debate. *Philosophy & Technology 36* (3).

Pifer, S. (2023). Russia, nuclear threats, and nuclear signaling. Technical report, The Brookings Institute. https://www.brookings.edu/articles/russia-nuclear-threats-and-nuclear-signaling/, Accessed: Oct. 17, 2023.

Scharre, P. (2023). *Four Battlegrounds*. New York, NY: W. W. Norton & Company.

Scharre, P. and K. Sayler (2016). Autonomous weapons and human control. Technical report, Center for a New American Security.

Sparrow, R. (2016). Robots and respect: Assessing the case against autonomous weapon systems. *Ethics & International Affairs 30* (01), 93–116.

Stop Killer Robots (2022). Negotiating a treaty on autonomous weapons systems: The way forward. Technical report, Stop Killer Robots.

Taddeo, M. and A. Blanchard (2022). A comparative analysis of the definitions of autonomous weapons systems. *Science and Engineering Ethics 28* (5), 1–22.

US Department of Defense (2023). DoD Directive 3000.09. Technical report, United States Department of Defense.

Williams, A. P. (2015). Defining autonomy in systems: Challenges and solutions. In A. P. Williams and P. D. Scharre (Eds.), *Autonomous Systems: Issues for Defense Policymakers*, pp. 27–62. The Hague, Netherlands: NATO Communications and Information Agency.

Winter, E. (2020). The compatibility of autonomous weapons with the principle of distinction in the law of armed conflict. *International & Comparative Law Quarterly 69* (4), 845–876.

Wong, Y. H., J. M. Yurchak, R. W. Button, A. Frank, B. Laird, O. A. Osoba,

R. Steeb, B. N. Harris, and S. J. Bae (2020). Deterrence in the age of thinking machines. Technical report, RAND Corporation, Santa Monica, CA.

Wood, N. G. (2023a). Autonomous weapon systems: A clarification. *Journal of Military Ethics 22*, 1–15.

Wood, N. G. (2023b). Autonomous weapon systems and responsibility gaps: A taxonomy. *Ethics and Information Technology 25* (1), 1–14.

Wood, N. G. (unpublished manuscript). Bombs, bots, and the principle of distinction. https://www.researchgate.net/publication/374913905.