Andrea Caligiuri (ed.), *Legal Technology Transformation. A Practical Assessment*
(Editoriale Scientifica 2020) **ISBN:** 978-88-9391-985-2

THE ETHICAL AND LEGAL IMPLICATIONS OF AUTONOMY IN WEAPONS SYSTEMS

DANIELE AMOROSO – GUGLIELMO TAMBURRINI

1. *Introduction*

According to the most accredited view, to count as autonomous, a weapons system must be able to select and engage targets without any human intervention after its activation.[1] Starting from this basic and quite inclusive condition, the Stockholm International Peace Research Institute (SIPRI) introduced additional distinctions between types of existing Autonomous Weapons Systems (AWS):[2] (i) *air defense systems* (e.g. Phalanx[3]); (ii) *active protection systems*, which shield armoured vehicles by identifying and intercepting anti-tank missiles and rockets (e.g. Trophy[4]); (iii) *robotic sentries*, like the Super aEgis II stationary robotic platform tasked with the surveillance of the demilitarized zone between North and South Korea;[5] (iv) *guided munitions*, like the Dual-Mode Brimstone, which are endowed with the capability of autonomously identifying and engaging targets that are not in sight of the attacking aircraft;[6] (v) *loitering munitions*, such as the Harpy NG,[7] which overfly an assigned area in search of targets to dive-bomb and destroy.

This classification stands in need of continual expansion on account of ongoing military research projects on unmanned ground, aerial and marine vehicles that are capable of autonomously performing targeting decisions. Notably, research work based on swarm intelligence technologies is paving the way to swarms of small-size and low-

[1] Crucial stakeholders have converged around this formulation of central properties of "autonomy": one of the technologically most advanced military powers (US Department of Defense 'Autonomy in Weapons Systems' Directive 3000.09 (21 November 2012)), the main international humanitarian organization (International Committee of the Red Cross (ICRC), 'Views on autonomous weapon system' Paper submitted to the Informal meeting of experts on lethal autonomous weapons systems of the Convention on Certain Conventional Weapons (Geneva, 11 April 2016)), and the coalition of NGOs advocating a ban on AWS (Campaign to Stop Killer Robots, 'Urgent Action Needed to Ban Fully Autonomous Weapons' Press release (23 April 2013)).

[2] V. Boulanin and M. Verbruggen, *Mapping the Development of Autonomy in Weapon Systems* (SIPRI 2017).

[3] R. H. Stoner, 'R2D2 with Attitude: The Story of the Phalanx Close-In Weapons' (30 October 2009) NavWeap <www.navweaps.com/index_tech/tech-103.htm>.

[4] 'Trophy Active Protection System' (10 April 2007) Defense Update <https://defense-update.com/20070410_trophy-2.html>

[5] S. Parkin, 'Killer Robots: The soldiers that never sleep' (16 July 2015) BBC Future <www.bbc.com/future/story/20150715-killer-robots-the-soldiers-that-never-sleep>.

[6] UK Royal Air Force, *Aircraft & Weapons* (2007) 87.

[7] D. Gettinger and A. H. Michel, *Loitering Munitions* (Center for the Study of the Drone 2017).

cost unmanned weapons systems. These are expected to overwhelm enemy defenses by their numbers and may additionally perform autonomously targeting functions.[8]

The technological realities and prospects of AWS raise a major ethical and legal issue: Is it is permissible to let a robotic system unleash destructive force and take attendant life-or-death decisions without any human intervention? Public awareness about the ethical and legal implications of autonomous targeting had been raised by the Campaign "Stop Killer Robots", which was launched in 2013 by an international coalition of NGOs with the primary goal of "ban[ning] lethal robot weapons".[9] Worldwide pressures from civil society prompted States to initiate discussion of normative frameworks to govern the design, development, deployment and use of AWS. Diplomatic dialogues on this topic have been conducted since 2014 at the United Nations in Geneva, within the institutional framework of the Convention on Certain Conventional Weapons (CCW). Informal Meetings of Experts on lethal autonomous weapons systems were held on an annual basis at the CCW in Geneva, from 2014 to 2016. Subsequently, the CCW created a Group of Governmental Experts (GGE) on lethal autonomous weapons systems (LAWS), which still remains (as of 2020) the main institutional forum where the issue of autonomy in weapons systems is annually debated at an international level. Various members of the robotics research community take part to the GGE's meetings.[10] So far, the main outcome of the GGE's work is the adoption by consensus of a non-binding instrument, that is, the 11 Guiding Principles on LAWS.[11]

An examination of the Guiding Principles, and more in general of the debate at the GGE, shows that discussions on AWS have been progressively zooming in on distinctive roles for the "human element" to play in the use of force. Indeed, it is widely agreed on by participants in these debates that the identification of normatively acceptable human-weapon interactions constitutes the keystone of any future regulation of AWS. Both the Campaign and an increasing number of States have come to maintain that a requirement should be established under international law, to rule that any weapons systems must be subject to meaningful human control (MHC).[12] It is exactly here, however, that international consensus stops. Indeed, it is far from settled – even among those favoring an MHC requirement – what its actual content should be or, to put it more sharply, what is normatively demanded to make human control over weapon systems truly "meaningful".

---

[8] See, among many others, P. Scharre, *Robotics on the Battlefield Part II. The Coming Swarm* (Center for a New American Security, October 2014) and M. Brehm and A. de Courcy Wheele, 'Swarms' Article36 discussion paper for the Convention on Certain Conventional Weapons (Geneva, March 2019).

[9] Campaign to Stop Killer Robots (n 1).

[10] More information on these meetings are available on the webpage of the UN Office at Geneva: <https://unog.ch>.

[11] High Contracting Parties to the Convention on Conventional Weapons, 'Final Report of the 2019 Meeting' (13 December 2019) UN Doc CCW/MSP/2019/CRP.2/Rev.1, Annex III ("Guiding Principles on Lethal AWS").

[12] The MHC formula made its first appearance in the AWS debate in a 2013 paper by the UK-based NGO Article 36 (*Killer Robots: UK Government Policy on Fully Autonomous Weapons* (Article 36, April 2013) <http://www.article36.org/wp-content/uploads/2013/04/Policy_Paper1.pdf>. For a survey of the States supporting the MHC requirement, see Human Rights Watch, *Stopping Killer Robots Country Positions on Banning Fully Autonomous Weapons and Retaining Human Control* (August 2020).

Against this background, this contribution is aimed to provide an overview of the AWS debate, with a focus on the MHC turning point and its ethical and legal underpinnings. In Section 2 a schematic account is provided of chief ethical and legal concerns about autonomy in weapons systems. Then, the main proposals regarding the MHC content are introduced and analyzed, including our own proposal of a "principled, differentiated and prudential" human control policy on AWS. Finally, it is pointed out how our proposal may help overcome the hurdles that are currently preventing the international community from adopting a legal regulation on the matter.

2. *Mapping the ethical and legal debate on AWS*

A clear outline of the main ethical and legal concerns raised by AWS is found already in a 2013 Report, significantly devoted to "lethal autonomous robotics and the protection of life", by the UN Special Rapporteur on extrajudicial, summary or arbitrary executions, Christof Heyns.[13] These concerns are profitably grouped under four headings: (i) *compliance with IHL*, (ii) *responsibility ascription problems*, (iii) *violations of human dignity*, and (iv) *increased risk for peace and international stability*. Let us briefly expand on each one of them, by reference to relevant sections in Heyns' report.

(i) Compliance with IHL would require capabilities that are presently possessed by humans only, and that no robot is likely to possess in the near future, i.e., to achieve situational awareness in unstructured warfare scenarios and to formulate appropriate judgments there (paras. 63-74);[14]
(ii) Autonomy in weapons systems would hinder responsibility ascriptions in case of wrongdoings, by removing human operators from the decision-making process (paras. 75-81);[15]
(iii) The deployment of lethal AWS would be an affront to human dignity, which dictates that decisions entailing human life deprivation should be reserved to humans (paras. 89-97);[16]
(iv) Autonomy in weapons systems would threaten in special ways international peace and stability, by making wars easier to wage on account of reduced numbers of involved soldiers, by laying the conditions for unpredictable

[13] 'Report of the Special Rapporteur on extrajudicial, summary or arbitrary executions' UN Doc. A/HRC/23/47, 9 April 2013
[14] For a critique to this argument, see M. N. Schmitt and J. S. Thurnher, '"Out of the Loop": Autonomous Weapon Systems and the Law of Armed Conflict' (2013) 4 Harvard National Security Journal 231; M. Sassòli, 'Autonomous Weapons and International Humanitarian Law: Advantages, Open Technical Questions and Legal Issues to be Clarified' (2014) 90 International Law Studies 308. See also, for a convincing rejoinder, T. Krupiy, 'Of Souls, Spirits and Ghosts: Transposing the Application of the Rules of Targeting to Lethal Autonomous Robots' (2015) 16 Melbourne Journal of International Law 145.
[15] For further discussion, see Carrie McDougall, 'Autonomous Weapon Systems and Accountability: Putting the Cart before the Horse' (2019) 20 Melbourne J of Intl L 58.
[16] For an overview of the debate on this issue, see Daniele Amoroso, *Autonomous Weapons Systems and International Law. A Study on Human-Machine Interactions in Ethically and Legally Sensitive Domains* (ESI/Nomos 2020), 161-215.

interactions between AWS and their harmful outcomes, by accelerating the pace of war beyond human reactive abilities (paras. 57-62).[17]

These sources of concern jointly make the case for claiming that a meaningful human control (MHC) over weapons systems should be retained exactly in the way of their critical target selection and engagement functions. Accordingly, the notion of MHC enters the debate on AWS as an ethically and legally motivated constraint on the use of any weapons systems, including autonomous ones. The issue of human-robot shared control in warfare is thereby addressed from a distinctive humanitarian perspective, insofar as autonomous targeting may impinge, and deeply so, upon the interests of persons and groups of persons that are worthy of protection from ethical or legal standpoints.

But what does MHC more precisely entail? What is normatively demanded to make human control over weapon systems truly "meaningful"? The current debate about AWS, which we now turn to consider, is chiefly aimed to provide an answer to these questions.

3. *Uniform policies for meaningful human control*

The foregoing ethical and legal reasons go a long way towards shaping the content of MHC, by pinpointing general functions that should be prescriptively assigned to humans in shared control regimes, and by providing general criteria to distinguish perfunctory from truly meaningful human control. More specifically, the ethical and legal reasons for MHC suggest a threefold role for human control on weapon systems to be "meaningful". First, the obligation to comply with IHL entails that human control must play the role of a *fail-safe actor*, contributing to prevent a malfunctioning of the weapon from resulting in a direct attack against the civilian population, or in excessive collateral damages.[18] Second, in order to avoid accountability gaps, human control is required to function as *accountability attractor*, i.e. to secure the legal conditions for responsibility ascription in case a weapon follows a course of action that is in breach of international law.[19] Third and finally, from the principle of human dignity respect it follows that human control should operate as a *moral agency enactor*, by ensuring that decisions affecting the life, physical integrity and property of people (including combatants) involved in armed conflicts are not taken by non-moral artificial agents.[20]

But how are human-weapon partnerships to be more precisely shaped on the basis of these broad constraints? Several attempts to answer this question have been made by parties involved in the AWS ethical and legal debate. The answers that we turn to examine now outline *uniform* human control policies, whereby one size of human

[17] On this issue, see J. Altmann and F. Sauer, 'Autonomous Weapon Systems and Strategic Stability' (2017) 59 Survival 117-142.

[18] P. Scharre, 'Centaur Warfighting: The False Choice of Humans v. Automation' (2016) 30 Temple International and Comparative Law Journal 151, 154.

[19] Th. Chengeta, 'Defining the emerging notion of "Meaningful Human Control" in autonomous weapon systems' (2017) 49 New York Journal of International Law and Politics 833, 888.

[20] ICRC, 'Ethics and autonomous weapon systems: An ethical basis for human control?' Working paper submitted to the Group of Governmental Experts on lethal autonomous weapons of the Convention on Conventional Weapons (Geneva, 29 March 2018) UN Doc CCW/GGE.1/2018/WP.5.

control is claimed to fit all AWS and each one of their possible uses. These are the "boxed autonomy", "denied autonomy" and "supervised autonomy" control policies.

The boxed autonomy policy assigns to humans the role of constraining the autonomy of a weapons system within an operational box, constituted by "predefined [target] parameters, a fixed time period and geographical borders".[21] Accordingly, the weapons system would be enabled to autonomously perform the critical functions of selecting and engaging targets, but only within the boundaries set forth by the human operator or the commander at the planning and activation stages.[22]

The boxed autonomy policy seems to befit a variety of *deliberate targeting* situations, which involve military objectives that human operators know in advance and can map with high confidence within a defined operational theatre. It seems, however, unsuitable to govern a variety of *dynamic targeting* situations. These require one to make changes on the fly to planned objectives and to pursue targets of opportunity. The latter are unknown to exist in advance (unanticipated targets) or else are not localizable in advance with sufficient precision in the operational area (unplanned targets). Under these conditions, boxed autonomy appears to be problematic from a normative perspective, insofar issues of distinction and proportionality that one cannot foresee at the activation stage may arise during mission execution.

By the same token, a boxed autonomy policy may not even suffice to govern deliberate targeting of military objectives placed in unstructured warfare scenarios. To illustrate, consider the loitering munition Harpy NG, which is endowed with the capability of patrolling for several hours a predefined box in search of enemy targets satisfying given parameters. The conditions licensing the activation of this loitering munition may become superseded if civilians enter the boxed area, erratic changes occur or surprise-seeking intentional behaviors are enacted.[23] Under these various circumstances, there is "fail-safe" work for human control to do at the mission execution stage too.

In sharp contrast with the boxed autonomy policy, the denied autonomy policy rules out any autonomy whatsoever for weapons systems in the critical targeting function, and therefore embodies a most restrictive interpretation of MHC.[24] Denied autonomy undoubtedly fulfils the threefold normative role for human control as fail-safe actor, accountability attractor, and moral agency enactor. However, this policy has been sensibly criticized for setting too high a threshold for machine autonomy, in ways that are divorced from "the reality of warfare and the weapons that have long been

[21] International Panel on the Regulation of Autonomous Weapons (IPRAW), *Focus on Human Control* (August 2019).

[22] This solution is advocated by the Dutch government, which endorsed on this point the recommendations issued by the Dutch Advisory Council on International Affairs (AIV) and Advisory Committee on Issues of Public International Law (CAVV) in their joint report 'Autonomous weapon systems: the need for meaningful human control' (2015) No. 97 AIV / No. 26 CAVV. For a more detailed account of this approach, see M. Roorda, 'NATO's Targeting Process: Ensuring Human Control Over and Lawful Use of 'Autonomous' Weapons', in A. Williams and P. Scharre (eds), *Autonomous Systems: Issues for Defence Policymakers* (NATO, 2015) 152; M. A. C. Ekelhof, 'Moving Beyond Semantics on Autonomous Weapons Systems: Meaningful Human Control in Operation' (2019) 10 Global Policy 343.

[23] D. Akerson, 'The Illegality of Offensive Lethal Autonomy', in D. Saxon (ed.), *International Humanitarian Law and the Changing Technology of War* (Brill/Nijhoff, 2013) 65, 87.

[24] Chengeta (n 19).

considered acceptable in conducting it."[25] To illustrate this criticism, consider air defensive systems, which autonomously detect, track, and target incoming projectiles. These systems have been aptly classified as SARMO weapons, where SARMO stands for "Sense and React to Military Objects". SARMO systems are hardly problematic from ethical and legal perspectives, in that "they are programmed to automatically perform a small set of defined actions repeatedly. They are used in highly structured and predictable environments that are relatively uncluttered with a very low risk of civilian harm. They are fixed base, even on Naval vessels, and have constant vigilant human evaluation and monitoring for rapid shutdown".[26]

SARMO systems expose the overly restrictive character of a denied autonomy policy. Thus, one wonders whether milder forms of human control might be equally able to strip the autonomy of weapons systems of its ethically and legally troubling implications. This is indeed the aim of the supervised autonomy policy, which occupies a middle ground between boxed and denied autonomy, insofar as it requires humans to be on-the-loop of AWS missions.

As defined in the US DoD Directive 3000.09 on "Autonomy in Weapons Systems", human-supervised AWS are designed "to provide human operators with the ability to intervene and terminate engagements, including in the event of a weapon system failure, before unacceptable levels of damage occur".[27] Notably, human-supervised AWS may be used for defending manned installations and platforms from "attempted time-critical or saturation attacks", provided that they do not select "humans as targets" (see, e.g., the Phalanx Close-In Weapons System in use on US surface combat ships).[28] While undoubtedly effective for these and other warfare scenarios, supervised autonomy is not the silver bullet for every ethical and legal concern raised by AWS. To begin with, by keeping humans on-the-loop one would not prevent faster and faster offensive AWS from being developed, eventually reducing the role of human operators to a perfunctory supervision of decisions taken at superhuman speed, while leaving the illusion that the human control requirement is still complied with.[29] Moreover, the automation bias – the human propensity to overtrust machine decision-making processes and outcomes – is demonstrably exacerbated by a distribution of control privileges that entrusts humans solely with the power of overriding decisions autonomously taken by the machines.[30]

To sum up. Each one of the boxed, denied, and supervised autonomy policies provides useful hints towards a normatively adequate human-machine shared control policy for military target selection and engagement. However, the complementary defects of these *uniform* control policies suggest the implausibility of solving the MHC problem with one formula, to be applied to all kinds of weapons systems and to each one of their possible uses. This point was consistently made by the US delegation at

[25] M. C. Horowitz and P. Scharre, *Meaningful Human Control in Weapon Systems: A Primer* (Center for a New American Security, March 2015).
[26] International Committee for Robot Arms Control, 'Statement on technical issues' delivered at the informal meeting of experts on lethal autonomous weapons (Geneva, 14 May 2014).
[27] US DoD Directive (n. 1), 13.
[28] Ibid., 3 para. 4(c)(2).
[29] E. Schwarz, 'The (Im)possibility of Meaningful Human Control for Lethal Autonomous Weapon Systems' (29 August 2018) Humanitarian Law & Policy <https://blogs.icrc.org/law-and-policy/2018/08/29/im-possibility-meaningful-human-control-lethal-autonomous-weapon-systems/>.
[30] L. J. Skitka, K. L. Mosier and M. Burdick, 'Does Automation Bias Decision-making?' (1999) 51 International Journal of Human-Computer Studies 991.

GGE meetings in Geneva: "there is not a fixed, one-size-fits-all level of human judgment that should be applied to every context".[31]

4. *The case for a prudential, differentiated and principled approach to meaningful human control*

Other approaches to MHC aim to reconcile the need for differentiated policies with the above ethical and legal constraints on human control. Differentiated policies modulate human control along various autonomy levels for weapons systems. A taxonomy of increasing autonomy levels concerning the AWS critical target selection and engagement functions was proposed by Noel Sharkey (and only slightly modified here, with regard to levels 4 and 5):[32]

L1. A human engages with and selects targets and initiates any attack.

L2. A program suggests alternative targets and a human chooses which to attack.

L3. A program selects targets, and a human must approve before the attack.

L4. A program selects and engages targets but is supervised by a human who retains the power to override its choices and abort the attack.

L5: A program selects targets and initiates attack on the basis of the mission goals as defined at the planning/activation stage, without further human involvement.

The main uniform control policies, including those examined in the previous section, are readily mapped onto one of these levels.

(L5) corresponds to the boxed autonomy policy, whereby MHC is exerted by human commanders at the planning stage of the targeting process only. As noted above, boxed autonomy does not constitute a sufficiently comprehensive and normatively acceptable form of human-machine shared control policy.

(L4) corresponds to the supervised autonomy policy. The uniform adoption of this level of human control must also be advised against, in the light of automation bias risks and increasing marginalization of human oversight. In certain operational conditions, however, it may constitute a normatively acceptable level of human control.

(L3) has been seldom discussed in the MHC debate. At this level, control privileges on critical targeting functions are equally distributed between weapon system (target selection) and human operator (target engagement). To the extent that the human deliberative role is limited to approving or rejecting targeting decisions suggested by the machine, this level of human control does not provide adequate bulwarks against the risk of automation bias.[33] In the same way as (L4), therefore, it should not be adopted as a general policy.

(L1) and (L2) correspond to shared control policies where the weapons system's autonomy is either totally absent (L1) or limited to the role of adviser and decision support system for human deliberation (L2). The adoption of these pervasive forms of

[31] United States, 'Human-Machine Interaction in the Development, Deployment and Use of Emerging Technologies in the Area of Lethal Autonomous Weapons Systems' Working paper submitted to the Group of Governmental Experts on lethal autonomous weapons of the Convention on Conventional Weapons (Geneva, 28 August 2018) UN Doc CCW/GGE.2/2018/WP.4.

[32] N. E. Sharkey, 'Staying the Loop: Human Supervisory Control of Weapons', in N. Bhuta et al. (eds), *Autonomous Weapons Systems. Law, Ethics, Policy* (CUP, 2016) 23, 34-37.

[33] M. L. Cummings, 'Automation and Accountability in Decision Support System Interface Design' (2006) The Journal of Technology Studies 23.

human control must also be advised against, insofar as some weapons (notably SARMO systems) have long been considered acceptable in warfare operations.

In the light of these difficulties, one might be tempted to conclude that the search for a comprehensive and normatively binding MHC policy should be given up, and that the best one can hope for is the exchange of good practices between States about AWS control, in addition to the proper application of national mechanisms to review the legality of weapons.[34] But alternatives are possible, which salvage the idea of a comprehensive MHC policy, without neglecting the need for differentiated levels of AWS autonomy in special cases. Indeed, the authors of this contribution have advanced the proposal of a comprehensive MHC policy, which is jointly *prudential*, *differentiated* and *principled*.[35]

The *prudential* character of this policy is embodied into the following default rule: high levels of human control L1-L2 should be exerted on all weapons systems and uses thereof, unless the latter are included in a list of exceptions agreed on by the international community of States. The prudential imposition by default of L1 and L2 is aimed at minimizing the risk of breaches of IHL, accountability gaps, or affronts to human dignity, should international consensus be lacking on whether, in relation to certain classes of weapons systems or uses thereof, higher levels of machine autonomy are equally able to grant the fulfilment of genuinely meaningful human control. The *differentiated* character of this policy is embodied in the possibility of introducing internationally agreed exceptions to the default rule. However, these exceptions should come with the indication of what level is required to ensure that the threefold role of MHC (fail-safe actor, accountability attractor, moral agency enactor) is adequately performed, which characterizes our approach as *principled*.

In the light of the above analysis, this should be done by taking into account at least the following observations:

a) The (L4) human supervision and veto level might be deemed as an acceptable level of control only in case of anti-materiel AWS with exclusively defensive functions (e.g. Phalanx or Iron Dome). In this case, ensuring that human operators have full control over every single targeting decision would pose a serious security risk, which makes the application of (L1), (L2), and (L3) problematic from both military and humanitarian perspectives. The same applies to active protection systems, like Trophy, provided that their use in supervised-autonomy mode is excluded in operational environments involving a high concentration of civilians.
b) (L1) and (L2) could also be impracticable in relation to certain missions because communication constraints would allow only limited bandwidth. In this case,

[34] See, e.g., the views expressed by Switzerland ('A "compliance-based" approach to Autonomous Weapon Systems' Working paper submitted to the Group of Governmental Experts on lethal autonomous weapons of the Convention on Conventional Weapons (Geneva, 10 November 2017) UN Doc CCW/GGE.1/2017/WP.9) and United States ('Statement for the General Exchange of Views' delivered at the Group of Governmental Experts on lethal autonomous weapons of the CCW (Geneva, 9 April 2018).

[35] For an early exposition of this approach, see Daniele Amoroso and Guglielmo Tamburrini. *What Makes Human Control over Weapon Systems "Meaningful"?* (International Committee for Robot Arms Control, August 2019). See also V. Boulanin et al., *Limits on Autonomy in Weapon Systems: Identifying Practical Elements of Human Control* (June 2020) SIPRI <https://www.sipri.org/sites/default/files/2020-06/2006_limits_of_autonomy_0.pdf>.

military considerations should be balanced against humanitarian ones. One might allow for less bandwidth-heavy (L3) control in two cases: deliberate targeting and dynamic targeting in fully structured scenarios, e.g. in high seas. In both hypotheses, indeed, the core targeting decisions have actually been taken by humans at the planning/activation stage. Unlike (L4), however, (L3) ensures that there is a human on the attacking end who can verify, in order to deny or grant approval, whether there have been changes in the battlespace which may affect the lawfulness of the operation. Looking at existing technologies, (L3) might be applied to sentry robots deployed in a fully structured environment, like the South-Korean Super aEgis II.

c) The (L5) boxed autonomy level should be considered incompatible with the MHC requirement, unless operational space and time frames are so strictly circumscribed to make targeting decisions entirely and reliably traceable to human operators.

5. *Concluding remarks*

Recent advances in autonomous military robotics have raised unprecedented ethical and legal issues. Regrettably, diplomatic discussions at the GGE in Geneva not only have so far fallen short of working out a veritable legal regime on meaningful human control over AWS, but – what is worse – are currently facing a stalemate, which is mainly determined by the opposition of major military powers, including the US and the Russian Federation, to the adoption of any kind of international regulation on the matter.[36]

Our proposal of relinquishing the quest for a one-size-fits-all solution to the MHC issue in favour of a suitably differentiated approach may help sidestep current stumbling blocks. Diplomatic and political discontent about an MHC requirement that is overly restrictive with respect to the limited autonomy of some weapons systems might indeed be mitigated recognising the possibility of negotiating exceptions to L1-L2 human control, by identifying weapons systems and contexts of use where milder forms of human control will suffice to ensure the fulfilment of the fail-safe, accountability, and moral agency properties whose preservation generally underpins the normative concerns about weapons' autonomy in targeting critical functions.

[36] D Lewis, 'An Enduring Impasse on Autonomous Weapons' (28 September 2020) Just Security <https://www.justsecurity.org/72610/an-enduring-impasse-on-autonomous-weapons>.