

## ONLINE ARTICLE

On the Pitfalls of Technophilic Reason: A Commentary on Kevin Jon Heller’s “The Concept of ‘the Human’ in the Critique of Autonomous Weapons”<sup>\*</sup>

---

*Elke Schwarz*<sup>†</sup>  
*Neil Renic*<sup>‡</sup>

---

<sup>\*</sup> This article is part of a symposium on Kevin Jon Heller’s “The Concept of “the Human” in the Critique of Autonomous Weapons,” published in this journal in 2023. All articles in the symposium can be found in the Harvard National Security Journal Online at <https://harvardnsj.org/onlineedition>.

<sup>†</sup> Elke Schwarz is Professor of Political Theory at Queen Mary University London. Her research focuses on the intersection of ethics of war and technology with an emphasis on unmanned and autonomous/intelligent military technologies and their impact on the politics of contemporary warfare. She is Vice-Chair for the International Committee for Robot Arms Control (ICRAC) and the author of *Death Machines: The Ethics of Violent Technologies*.

<sup>‡</sup> Neil Renic is a lecturer in ethics at the University of New South Wales, a fellow at the Centre for Military Studies at the University of Copenhagen, and member of the International Committee for Robot Arms Control (ICRAC). He specialises in the changing character of war, the ethics of killing, and emerging military technologies.

## Table of Contents

<b>I. INTRODUCTION .....</b>	<b>2</b>
<b>II. PITFALLS IN HELLER’S CRITIQUE OF THE AWS CRITICS .....</b>	<b>2</b>
<i>A. Even If Machines Do Not Make Decisions Like A Human Does, They Deserve Scrutiny For Transforming The Role Of Human Decision-Makers. ....</i>	<i>3</i>
<i>B. AWS Have Greater Autonomy From Operator Intentions Than Prior Weapon Systems .....</i>	<i>3</i>
<i>C. The Actions Of AWS Cannot Be Held Accountable By Laws Intended To Regulate Morally Meaningful Human Action .....</i>	<i>4</i>
<i>D. Warfare Cannot Be Judged Solely By Standards Of Technical Effectiveness .....</i>	<i>5</i>
<b>III. THE PURPOSE OF AN AUTONOMOUS WEAPON SYSTEM IS WHAT IT DOES .....</b>	<b>6</b>
<b>IV. CONCLUSION .....</b>	<b>9</b>

## I. INTRODUCTION

Critiquing a critique is a delicate matter. One risk is that the intention of the critique (either one) is distorted or misconstrued. The motivations of the respective authors can be difficult to grasp with precision, and we all read texts and ideas through the lens of our own experiences and perspectives, leading, invariably, to different, often incompatible interpretations.

Thankfully, this difficult business is made easier by the fact that we know and like Kevin (referred to as “Heller” from hereon). Our disagreements with Heller over the potential benefits and dangers of autonomous weapons – or autonomous weapons systems (AWS) – are substantial. Yet both his and our approach ultimately stem from the same place: a concern with the human in warfare, a commitment to battlefield restraint, and a drive to ensure that armed conflict is waged less frequently and less cruelly. From these shared foundations, healthy debate can be had.

Even those who contest Heller’s stance on autonomous weapons can appreciate his breadth of analysis. His article provides several critical reminders.<sup>1</sup> First, it offers a reminder to avoid anthropomorphizing the autonomous weapon. For both proponents and critics of this technology, the human must remain key. Second, it argues in favor of taking history seriously. Not every challenge posed by autonomous weapons is unprecedented. Third, it engages the human agents of war as they are (cognitively and morally flawed), rather than as we wish them to be. And lastly, it provides a reminder that “humans on-the-loop” – frequently framed as a middle ground compromise between proponents and critics of this technology – is no panacea to the moral and legal challenges of autonomous weapons. Heller offers an intellectually clear and robust defense of autonomous weapons as a (potential) moral and legal improvement over human combatants in war. For those who disagree, as we most certainly do, his work forces careful reflection on the nature of this disagreement.

The key points raised in our response to Heller are three-fold. First, we begin with a brief discussion on the pitfalls of critiquing AWS critics through a techno-centric lens. The second section addresses Heller’s juxtaposition of humans and AWS with respect to IHL compliance. Steering away from the realm of abstraction, we situate the use of AWS in the real world. While humans will always play a role within these systems, and thus war itself, the role and agency of such humans is likely to change. We argue that warfare within this human-AWS nexus is likely to be less “humane,” not more. In a slightly unorthodox manner, we conclude by raising a question. We ask to what degree it is fruitful to approach this issue (as Heller mostly does) through legal abstraction, and reflect on the responsibility we have as scholars and researchers for the ideas and arguments we set loose in the world.

## II. PITFALLS IN HELLER’S CRITIQUE OF THE AWS CRITICS

Heller unpersuasively criticizes the deontological arguments made by AWS skeptics. He opens the discussion with an indicative inversion. Where critics claim, on deontological grounds, that autonomous weapons are *mala in se* (meaning evil in themselves or intrinsically bad) based on their characteristics and their externality to IHL,<sup>2</sup> Heller provocatively titles his section on deontological arguments “Humans as *bonum in se*” (as good in themselves, or intrinsically good). He then proceeds to rebut a range of arguments that cast humans as intrinsic

---

<sup>1</sup> Kevin Jon Heller, *The Concept of “The Human” in the Critique of Autonomous Weapons*, 15 HARV. NAT’L SEC. J. 1, 6 (2023).

<sup>2</sup> See generally Rob Sparrow, *Robots and Respect: Assessing the Case Against Autonomous Weapon Systems*, 30 ETHICS & INT’L AFFS. 93 (2016); Christian Nicholas Braun, *LAWS and the Mala in Se Argument*, 33 PEACE REV. 237 (2022).

bearers of certain qualities that are morally relevant to the business of killing. He lists these qualities in summary as follows: “only humans have morality”;<sup>3</sup> “only humans suffer”;<sup>4</sup> “only humans risk”;<sup>5</sup> “and only humans can kill with dignity.”<sup>6</sup>

We find the charges Heller raises against these specific, selectively chosen arguments to be overstated, missing the ethos of the more carefully crafted arguments put forward by AWS critics. No criticism of AWS (as far as we are aware) *hinges* on an understanding of humans in war as intrinsically, or even mostly good. Few, if any, opponents of this technology argue that killing can only ever be legitimate if the agent of violence physically risks his or her own life.<sup>7</sup> Critics understand that weaponized autonomy and “free will” are different, and that criticism of the former need not (and should not) presume the existence of the latter.

The deontological criticisms Heller rebuts are subtler than he credits. They categorically reject the morality of autonomous weapons use without indulging in a romanticised ideal of either humans or war. The inhumanity of AWS, to be sure, *does* matter, even if humans are not ideal actors. It matters that these machines are not human; that they cannot understand human contexts; are unable to draw on plural experiences of their own; are not sensitive to vulnerability; and lack a conception of the value of life, or indeed, the horrors of a violent death. Critics are right to doubt whether autonomous weapons can align with IHL frameworks, designed as these frameworks are with human capacities and limitations in mind. Their argument is one grounded in *compatibility* between flawed humans and flawed law. It does not require unjustified faith in the unassailable goodness of “humanity.”

Heller’s responses to deontological AWS criticisms understate the ways in which machines replace human decision-making and why this replacement matters.

*A. Even If Machines Do Not Make Decisions Like a Human Does, They Deserve Scrutiny for Transforming the Role of Human Decision-Makers*

For Heller, the term “decision making” is misused as applied to AWS. He argues that “the ‘moral judgment’ objection to machine killing necessarily assumes that an autonomous weapon ‘decides’ to take human life in a manner akin to human decision-making.”<sup>8</sup> In his view, this unduly anthropomorphizes the machine. However, the claim that the moral weight of killing can only be experienced by a human decision maker is not necessarily the same as claiming that the machine makes decisions as a human would. At issue is a machine replacing a role that a human inhabits in the wider configuration of technologized warfare, a role that was hitherto understood to rest on human judgement. It is the moral weight of *this* change in the character of war that warrants scrutiny.

*B. AWS Have Greater Autonomy from Operator Intentions than Prior Weapon Systems*

The alleged anthropomorphizing of AWS, Heller goes on to argue, occludes the fact that a human programs the technology and is “responsible for determining which kinds of individuals and objects the machine will target.”<sup>9</sup> Understood in this sense, the AWS is a mere tool of human will, actualizing operator intentions. But can this argument really stand as a rebuttal? Is

---

<sup>3</sup> Heller, *supra* note 1, at 6.

<sup>4</sup> *Id.* at 11.

<sup>5</sup> *Id.* at 13.

<sup>6</sup> *Id.* at 14.

<sup>7</sup> See generally NEIL RENIC, ASYMMETRIC KILLING: RISK AVOIDANCE, JUST WAR, AND THE WARRIOR ETHOS (2020).

<sup>8</sup> Heller, *supra* note 1, at 7.

<sup>9</sup> *Id.* at 8.

pre-programming a system, whose specific actions cannot be known with exactitude (as Heller himself notes, “it is impossible to know *ex ante* each and every target it will engage”),<sup>10</sup> really the same as making and affecting the in-the-moment decision to kill? Is the human agency and judgement we consider to be morally relevant for taking another human life present in sufficient quantities in both instances?

In Heller’s view, the answer is yes. Here, he references a number of antecedent weapons and practices, such as The Mark 46 and the CAPTOR, which have a similar, although not equal, distance between human operator intent and effect.<sup>11</sup> But the CAPTOR and the Mark 46 are not comparable to the AI-enabled weapon systems that most concern critics. The level of autonomy and dynamic “decision-making” differs significantly. The systems that cause the greatest concern will employ AI for *target identification* based on a much broader set of parameters and the execution of the targeting function will be done with a significantly higher level of autonomy.<sup>12</sup>

To further argue that AWS does not represent a novel move away from human agency, Heller offers up the practice of “reconnaissance by fire,” in which targets are identified according to space, not individual characteristics. He uses this to illustrate that in certain warfare practices, there is a functional distance between the broader objective and the soldiers’ knowledge of what targets they are firing upon.<sup>13</sup> But knowledge about target characteristics is not the same as understanding what a human target is and what it means to take it out. Even without knowing who is being killed, a soldier’s understanding of the decision to kill remains complex. It is more than an automatic execution of targeting parameters, and more than the functional fulfillment of a standing order, no matter how routinized and mechanized warfighting becomes. Even if we wish to distill the application of violence down to the simple execution of targeting rules, this has hitherto been coupled with an assumption of meaningful human control. In other words, existing warfare practices like “reconnaissance by fire” presume control by a human who *understands* and can take responsibility (in the broadest sense) for the act of violence. An autonomous weapon system does not have this capacity.

The antecedent technologies and practices Heller references should not be ignored. At the very least, they force critics of AWS to reflect on important questions: what, if anything, is new, and if new, what, if anything, is problematic? The history of war is important in these discussions, not because it confirms autonomous violence as an entirely novel challenge (it is not), but rather as a means through which to better appreciate the moral and legal challenges replicated, and in many cases, intensified by this technology.<sup>14</sup>

### *C. The Actions of AWS Cannot Be Held Accountable by Laws Intended to Regulate Morally Meaningful Human Action*

Critics of AWS take issue with the inability of these systems to bear the moral weight of the kill decision. Such critics, do not, as Heller claims, regard autonomous weapons as “a subject capable of morally meaningful action.”<sup>15</sup> The act of killing is morally meaningful for humans only. A machine cannot understand this meaning. In the absence of such

---

<sup>10</sup> *Id.* at 9.

<sup>11</sup> *Id.* at 9–10.

<sup>12</sup> See generally Elke Schwarz, *Autonomous Weapon Systems, Artificial Intelligence and the Problem of Meaningful Human Control*, 5 THE PHIL. J. OF CONFLICT AND VIOLENCE 53 (2021).

<sup>13</sup> *Id.* at 10.

<sup>14</sup> See generally Neil Renic & Elke Schwarz, *Crimes of Dispassion: Autonomous Weapons and the Moral Challenge of Systematic Killing*, 37 ETHICS & INT’L AFFS. 321 (2023).

<sup>15</sup> Heller, *supra* note 1, at 9 (quoting PAUL SCHARRE & MICHAEL C. HOROWITZ, AN INTRODUCTION TO AUTONOMY IN WEAPON SYSTEMS 16 (2016)).

understanding, machines must not be empowered with the human decision to identify specific targets from an often over-broad set of parameters.

Far from being ignored, the ontological and epistemological distance between a machine “decision” and a human decision is significant to critics of AWS. Ironically, perhaps, Heller and his interlocutors might agree that machines do not make decisions in the human sense. But this insight has very different implications for each party. For Heller, this is a functional difference that suggests the utilitarian advantages of AWS. Critics, by contrast, raise deontological objections, suggesting that it is precisely because we cannot equate machines and humans that judging them by the same effectiveness standards makes no sense. This unbridgeable gap renders autonomous machines intrinsically evil as instruments of killing, permanently outside the human institution of law. Although we cannot speak for the various critics Heller engages in his discussion, our interpretation of the arguments advanced by Krishnan, Rosert and Sauer, Heyns, and others is that human life makes sense only in the context of human relations.<sup>16</sup> Legal frameworks, including IHL, and much of the Just War Tradition, have been formed specifically around the assumption of human agency, an assumption that includes the understanding that humans are flawed and fallible.<sup>17</sup> Humans are not machines and human affairs cannot and should not be read exclusively through the narrow lens of technological capacity.

The question then is one that many critics of autonomous weapons have picked up: what is required to adhere to the rules and standards that govern us? This is a philosophical question, as well as a practical one. It is a question of moral, not merely technical, significance, and should remain so, regardless of how mechanized and systematized contemporary warfare has become.<sup>18</sup> Answering the question involves rendering human judgment on and taking human responsibility for life and death. Certain technological capabilities, including autonomous weapons, complicate the attainment of both judgment and responsibility. We should not mistake human agency for technological agency, lest we wish to court disaster.<sup>19</sup>

#### *D. Warfare Cannot Be Judged Solely by Standards of Technical Effectiveness*

In Heller’s essay, the deck is stacked in favour of the technological. He closes the first section with an indicative assessment:

The most fundamental problem with deontological objections, however, is precisely their deontology. Because deontological arguments are by definition non-consequentialist, they would prohibit states from using autonomous weapons in conflict *even if doing so would lead to fewer civilian casualties and less unnecessary combat suffering*.<sup>20</sup>

With this, Heller rejects the deontological premise that ethical considerations raised on grounds of duty or dignity, or an incompatibility between machine processes and the ethos of

---

<sup>16</sup> See ARMIN KRISHNAN, KILLER ROBOTS: LEGALITY AND ETHICALITY OF AUTONOMOUS WEAPONS (2010); Elvira Rosert & Frank Sauer, *Prohibiting Autonomous Weapons: Put Human Dignity First*, 10 GLOB. POL’Y 370, 370 (2019); Christof Heyns, *Autonomous Weapons Systems: Living a Dignified Life and Dying a Dignified Death*, in AUTONOMOUS WEAPONS SYSTEMS: LAW, ETHICS, POLICY 3, 11 (Nehal Bhuta et al. eds., 2016).

<sup>17</sup> Christof Heyns, *Autonomous Weapons in Armed Conflict and the Right to a Dignified Life: An African Perspective*, 33 S. AFR. J. ON HUM. RTS., 51–52 (2017)

<sup>18</sup> Norbert Wiener knew this at the dawn of automated and autonomous weapons technology. See Norbert Wiener, *Some Moral and Technical Consequences of Automation*, 131 SCIENCE 1355, 1358 (1960).

<sup>19</sup> See *id.*

<sup>20</sup> Heller, *supra* note 1, at 17.

human-oriented legal parameters, have sufficient validity. It is also assumed that the costs and benefits of AWS could straightforwardly be ascertained, presumably with techniques adequate for that purpose. This perspective betrays a scientific-technological lens increasingly prominent in abstract reasoning.<sup>21</sup> The lens truncates the “experiential and situated knowledge” so important in warfare,<sup>22</sup> instead prioritising that which can be counted, measured, and assessed in an overly optimistic techno-solutionist spirit.

The charge that autonomous weapons critics unduly anthropomorphize the technology could be inverted in relation to Heller’s stance: he unduly technologizes the human in the process of killing. And this machinist logic weaves through the rest of the essay. In discussing the many challenges humans present to perfect compliance with IHL, Heller notes a series of biases: cognitive bias, availability bias, imaginability bias, base-rate bias, anchoring bias, overconfidence bias, object use bias, confirmation bias and stereotyping. He suggests, rightly, that these individual psychological traits make it difficult for the human to achieve compliance with IHL. In a section on “debiasing,” he explains the difficulty in ameliorating these flaws: “[t]raining people in statistical reasoning has had some success for simple tasks, but it ‘has not typically been fully tested in complex environments using unfamiliar and abstract rules.’”<sup>23</sup> One might be forgiven for momentary confusion: is this a discussion of human capabilities, or the capabilities of a robotic system? The discussion he offers presses human capabilities into an analytical framework fit for the latter and analyzes warfare solely as an engineering problem.

Imposing this mechanistic model on human behaviour for a side-by-side comparison is not productive towards furthering the pressing moral and legal debate on autonomy in weapon systems. Elsewhere, Heller states that “some of [the principle of distinction’s] ... central requirements, such as identifying combatants and recognising surrender, *involve little more than object recognition*.”<sup>24</sup> This narrow conception of IHL, and the role of the human within it, once again stacks the odds in favor of machine logics.

Heller’s excessively technologized approach pathologizes the human agent in war as not merely flawed, but irredeemably so. The same approach also gives the autonomous weaponry tasked with our replacement too easy a ride.

### III. THE PURPOSE OF AN AUTONOMOUS WEAPON SYSTEM IS WHAT IT DOES

Even if warfare were judged solely according to consequentialist standards for effectiveness, AWS would still be a flawed solution. Much of Heller’s essay is dedicated to exposing the failings of humans in war. Measured against a set of technical standards, the irrational human is found wanting, permanently doomed to lack the necessary capacity and inclination to functionally adhere to IHL. Viewed in such terms by Heller, the age-old question of how to tame the destructive excesses of war has an obvious answer:

[G]iven the irrationality of human decision-making, particularly in combat, warfare that eliminates human control – that is fought autonomously – promises *less* unnecessary death and destruction, not *more*. Put more simply: in terms of IHL compliance, humans are the problem, not the solution.<sup>25</sup>

<sup>21</sup> See generally Elke Schwarz, *Technology and Moral Vacuums in Just War Theorising*, 14 J. INT’L POL. THEORY 280 (2018); Elke Schwarz, *Trolleyology: Algorithmic Ethics for Killer Robots*, in HANDBOOK ON THE ETHICS OF ARTIFICIAL INTELLIGENCE (David Gunkel ed., 2024).

<sup>22</sup> See generally CHRISTOPHER COKER, *ETHICS AND WAR IN THE 21<sup>ST</sup> CENTURY* (2008).

<sup>23</sup> Heller, *supra* note 1, at 48.

<sup>24</sup> *Id.* at 31 (emphasis added).

<sup>25</sup> *Id.* at 54 (emphasis in original).

We disagree. “The purpose of a system is what it does” is a systems thinking heuristic first coined by Stafford Beer, who argues that there is “no point in claiming that the purpose of a system is to do what it constantly fails to do.”<sup>26</sup> This applies equally to autonomous weapons systems. Heller’s position, we argue, cannot be sustained once the human is compared, as it must be, to autonomous violence as it *actually* manifests, rather than to an imagined system of perfect rationality.

Pioneering programs that combine sophisticated and dynamic target nomination systems with weapons platforms, such as the U.S. developed Project Maven, have the potential for fully autonomous use.<sup>27</sup> These systems are specifically designed not to merely identify a torpedo or a tank, but to be used in a dynamic process. They sift through a massive volume of data to nominate and execute targets autonomously at an accelerated pace. At this stage, a human remains in the loop to sign off on suggested targets, but this sign off also becomes a mechanized task where “decision-makers” are empowered, and required, to authorise “as many as 80 targets in an hour of work” — a “rapid staccato of ‘Accept, Accept, Accept.’”<sup>28</sup> Moreover, AWS of this nature are typically connected to multiple AI systems which work in layers when performing targeting tasks. Speed and volume are clearly prioritized as key to greater lethality.

In order to take the critiques of AWS seriously, one must take the technologies the critics critique seriously. Much hinges on a system’s actual capabilities, not its future promise, which is often wildly overstated and overhyped. Systems that employ AI for the full kill-chain are likely to be marred by incomplete, low-quality, incorrect, or discrepant data.<sup>29</sup> This, in turn, will lead to highly brittle systems and biased, harmful outcomes. Autonomous systems tend to be built and tested on rather limited samples of data. Sometimes it is synthetic data, and sometimes it is inappropriate data.<sup>30</sup> This is problematic enough to question the use of AWS, even before considering the messy complexities of the battlefield.

In addition to this, research in machine learning for AI is nowhere close to as scientifically robust as it might appear from media accounts. We know this, for example, from the healthcare industry, where AI systems have had mixed results in applications ranging from health predictions to diagnostics, often on account of overpromising the utility of AI for these tasks.<sup>31</sup> AI has value primarily for quite clearly delineated, narrow tasks, executed in closed systemic environments. Applications of AI elsewhere are highly likely to lead to faulty decisions, misapplications of force, and other harmful errors. System outcomes are inherently unpredictable, and the probabilistic nature of AI reasoning implicitly recognizes error and accident as a feature, not a bug, of the system.

---

<sup>26</sup> David Komlos & David Benjamin, *The Purpose of a System Is What it Does Not What It Claims To Do*, FORBES (Sept. 13, 2021), <https://web.archive.org/web/20210913130445/https://www.forbes.com/sites/benjaminkomlos/2021/09/13/the-purpose-of-a-system-is-what-it-does-not-what-it-claims-to-do/>.

<sup>27</sup> Katrina Manson, *AI Warfare is Already Here*, BLOOMBERG BUSINESSWEEK (Feb. 28, 2024), <https://www.bloomberg.com/features/2024-ai-warfare-project-maven/>.

<sup>28</sup> *Id.*

<sup>29</sup> See ARTHUR HOLLAND MICHEL, KNOWN UNKNOWN: DATA ISSUES AND MILITARY AUTONOMOUS SYSTEMS 3–5 (2021), <https://unidir.org/publication/known-unknowns-data-issues-and-military-autonomous-systems/>.

<sup>30</sup> *Id.*

<sup>31</sup> See generally Hadyia Faheem & Sanjib Dutta, *Artificial Intelligence Failure at IBM ‘Watson for Oncology’*, 21 IUP J. KNOWLEDGE MGMT. 47 (2023); David Kampmann, *Venture Capital, the Fetish of Artificial Intelligence and the Contradictions of Making Intangible Assets*, 53 ECON. SOC. 39 (2024); Eliza Strickland, *IBM Watson, Heal Thyself: How IBM Overpromised and Underdelivered on AI Health Care*, 56 IEEE SPECTRUM 24 (2019).



Research in computer vision is also not nearly as robust as Heller suggests.<sup>32</sup> In fact, computer vision faces a number of unresolved challenges, especially when it comes to human action recognition. This includes issues around the detection of variations among human features, “image-quality and frame rate” issues, “multiview variations” especially in uncontrolled settings, and general environmental issues, like poor visibility through bad weather or other background “noise” that may distort the relevant input data for an accurate and reliable result.<sup>33</sup> Detecting violent action is particularly challenging “because the available violence dataset is insufficient for deep network training. Also, human behaviour contains high intra-class variations and inter-class similarities that make violence detection very challenging.”<sup>34</sup> Some of these issues are likely to be mitigated as more data becomes available, hardware becomes more sophisticated, and technology generally advances. Others, critically, may not. We are dealing with an *open question*, a fact that should not be forgotten in these debates. The insistence that autonomous weapons will morally outperform humans in battle, if not now then one day, should be seen for what it is: an article of faith.

As empirics on machine learning, computer vision, and the actual use of AWS reinforces, the perfectibility of this technology is a highly speculative assumption, not a given. This is important to repeat, given the determination of many AI developers to convince states and the public of the opposite. Here, it is worth further interrogating the examples provided by Heller to illustrate his confidence in such technology. Quoting Elliot Winter, Heller writes that “not only can machines ‘observe at least as well as humans and, indeed, at higher resolution and with greater rapidity,’ their recognition ability ‘has now advanced to a point where it has reached parity with human recognition abilities.’”<sup>35</sup> Heller then substantiates this with reference to the company material of Malong Technologies, which boasts of a 94.78% accuracy in object recognition.<sup>36</sup> Elsewhere, Heller asserts that autonomous weapons systems will likely outmatch humans in differentiating between real and fake guns, an important ability in war, especially in complex urban operations. For evidence, Heller draws on the Patriot One’s “PatScan” threat-detection product, a system able to identify weapons with a “nearly 95% certainty.”<sup>37</sup> This statistic stems from a company promotional video for a product designed specifically for use within a civilian context to recognize concealed weapons in shopping malls or hotels.

These industry claims deserve at least as much scrutiny as Heller gives the human. AI discourse is awash with exaggerated claims regarding the efficacy of such systems, pushed forward by private actors with a direct financial interest in normalizing the technology and lowering barriers for use. Even when the technical claims provided by companies are accurate, they are benchmarked and assessed in very limited settings.<sup>38</sup> Too often, these limitations are ignored or downplayed, with specific data about specific technologies uncritically extrapolated and misapplied to the radically different war context.

---

<sup>32</sup> Heller, *supra* note 1, at 20.

<sup>33</sup> See Imen Jegham et al., *Vision-Based Human Action Recognition: An Overview and Real World Challenges*, 32 FORENSIC SCI. INT’L: DIGIT. INVESTIGATION 1, 2 (2020).

<sup>34</sup> See Tahereh Zarrat Ehsan et al., *An Accurate Violence Detection Framework Using Unsupervised Spatial–Temporal Action Translation Network*, 40 THE VISUAL COMPUT. 1515, 1515 (2024).

<sup>35</sup> Heller, *supra* note 1, at 21 (quoting Elliot Winter, *The Compatibility of Autonomous Weapons with the Principle of Distinction in the Law of Armed Conflict*, 69 ICLQ 845, 859, 867 (2020)).

<sup>36</sup> *Id.*

<sup>37</sup> *Id.* at 26.

<sup>38</sup> See Arthur Holland Michel, *Recalibrating Assumptions on AI: Towards an Evidence-Based and Inclusive AI Policy Discourse* 30, CHATHAM HOUSE (Apr. 12, 2023), <https://doi.org/10.55317/9781784135621>.

When it comes to the battlefield implementation of AI and machine learning, a high degree of caution is needed. As Paddy Walker notes, “[r]esearch demonstrates that the introduction of new parameters or slightly heterogeneous data to the data under which the weapon has been trained will confound autonomous weapon ML processes for the task envisaged, particularly in the changing nature of a battlespace with its prevalence of hidden, partially observable or camouflaged states.”<sup>39</sup> Autonomous weapon systems will need to be *trained*, and not merely programmed – any “marginally different set up, discrepancies between training dataset and sensed information or data poisoning arising from adversarial intervention may all lead to substantial variation” in the system’s action.<sup>40</sup> Add to this the “technical debt [that] arises from inappropriate architecture, shortcuts resulting from commercial pressures, poor testing protocols and poor holistic understanding of LAWS,”<sup>41</sup> and the idealisation of AWS in Heller’s piece becomes less and less tenable.

It is not our intention to deny the very real and enduring problems with human operators in war, nor do we wish to suppress discussion of more humane alternatives to the military status-quo. But humility is needed. The drive to humanize war must begin with a recognition of how resistant many of the moral problems of the battlefield ultimately are to technical solution. We also need to avoid an uncritical embrace of industry claims that cynically or genuinely assume away the immutable complexities of war.

As depressing an admission as it is, war can always get worse. We differ fundamentally from Heller, in recognizing the potential for autonomous weapons to make warfare even more violent, brutal, and unstable. This view is grounded not in a quixotic and romanticised view of the human, but rather a deep and justified caution over the alleged promise of autonomous weapons.

#### IV. CONCLUSION

By way of concluding, we would like to raise a question on the potential risks of autonomous weapons advocacy. Heller’s concern, like ours, is with ensuring that war is neither commenced nor conducted unjustly. But does argumentation and advocacy in the mode of Heller ultimately help our efforts to stand against impulses of war? Our answer is no. Claiming, contrary to the interpretation and intuition of most, that “[human] understanding is far less necessary to IHL than AWS critics assume,”<sup>42</sup> is to abstract the letter of the law from the ethos. Heller seeks to provide a fair and, as he sees it, overdue critique of the human in war, and he encourages consideration of autonomous weapons as a potential moral and legal improvement. What he actually does, we fear, is wrongly elevate this technology to a status it is not close to deserving.

There are certainly grounds for pessimism about human qualities in warfare, but Heller’s own pessimism in the human is supplemented by a somewhat unfounded optimism in technology and its perfectibility. He stacks the odds, placing the human agent into direct competition with an ideal machine entity that “perceive[s] the world accurately, understand[s] rationally, quarantines negative emotions and reliably translates thought into action.”<sup>43</sup> Missing is necessary reflection on the serious limits and dangers of such technological systems, and the functional embedding of humans within them. Also missing is an exploration of the

---

<sup>39</sup> Paddy Walker, *Leadership Challenges from the Deployment of Lethal Autonomous Weapon Systems*, 133 *RUSI JOURNAL* 10, 14 (2021).

<sup>40</sup> *Id.* at 15.

<sup>41</sup> *Id.* at 19.

<sup>42</sup> Heller, *supra* note 1, at 5.

<sup>43</sup> *Id.* at 4.

assumptions and motivations of those who most forcefully push for this technology. There is a fetishization of speed underpinning both the design and use of AWS, a “move fast and break things” logic driving innovation, and a spectre of great power competition invoked to silence criticism from those justifiably concerned with both the pace and character of this change. These factors matter immensely when considering how this technology will be used.

We need to think very seriously about the modes of war we should and should not countenance going forward, and the critical role of technology in fulfilling these visions. Morally and legally improving the battlefield is frustrating work. Positive change, when achievable at all, is rarely satisfying – too slow, too partial, too contingent. But beware the lure of the technological shortcut. History provides us with endless examples of human weakness and human cruelty in war. History also illustrates the acute moral danger when individual human agency is stripped away or surrendered to excessively systematic and anonymizing processes of violence. Such conditions open the way to battlefield misconduct or worse. Humanity is an endless disappointment in war, but we suspect that we will miss it when it is gone, especially if it is cleared away to make room for flawed technologies that fall short of their promise and deaden our moral imagination.