

Crimes of Dispassion: Autonomous Weapons and the Moral Challenge of Systematic Killing

Neil Renic  and Elke Schwarz¹ 

In June 1959, the German philosopher Günther Anders penned a letter to Claude Eatherly. Eatherly was a former U.S. Air Force pilot and then psychiatric patient, who experienced immense guilt over his relatively minor role in the Hiroshima bombing. In the correspondence that followed, both men wrote in detail about the event and their concern over what they saw as a gulf between the moral imagination of humanity and the material destructiveness of the new atomic age. Anders feared the “technification’ of our being” and the loss of agency this would invariably entail:

The fact that to-day it is possible that unknowingly and indirectly, like screws in a machine, we can be used in actions, the effects of which are beyond the horizon of our eyes and imagination, and of which, could we imagine them, we could not approve—this fact has changed the very foundations of our moral existence.²

For Anders, machine-logics were a potential—and potentially fatal—threat to conscience.³ In his 1950 book *The Human Use of Human Beings: Cybernetics and Society*, American mathematician and cybernetics pioneer Norbert Wiener similarly noted, with apprehension, that when humans are “knit into an organization in which they are used, not in their full right as responsible human beings, but as cogs and levers and rods, it matters little that their raw material is flesh and blood.

Neil Renic, Centre for Military Studies, University of Copenhagen, Copenhagen, Denmark (neil.renic@ifs.ku.dk)

Elke Schwarz, Queen Mary University of London, London, England (e.schwarz@qmul.ac.uk)

Ethics & International Affairs, 37, no. 3 (2023), pp. 321–343.

© The Author(s), 2023. Published by Cambridge University Press on behalf of Carnegie Council for Ethics in International Affairs. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

doi:10.1017/S0892679423000291

What is used as an element in a machine, is in fact an element in the machine."⁴ The concern for Wiener, as for Anders, was that the increased tendency toward human technification (the substitution of technology for human labor) and systematization would exacerbate the dispassionate application of lethal force and lead to more, not less, violence.

This insight is as apposite today as it was then as we face a future of accelerated and increasingly autonomous modes of highly systematized warfare. In particular, the scale and speed of the rollout of AI-enabled weapons systems should prompt reflection on the moral implications of this integration of nonhuman logics and systems into existing processes of military violence.

Systems are near omnipresent in any task requiring concerted human effort. Critically though, there are limits on the type and degree of systematization that are appropriate in human conduct, especially when it comes to collective violence. The systematic application of violence has been a feature of some of the most destructive episodes in modern human history, including colonial warfare, ethnic cleansing, and genocide. While each of these episodes is unique, commonalities can be identified among the processes: targeted peoples are classified by certain characteristics and organized into a pathologized category; violence is applied instrumentally and often dispassionately via systems of diffused responsibility; and the killing is in tension with moral values. Engaging these antecedents, we draw out the parallels in process between historical episodes of systematic violence and lethal autonomous weapon systems, a mode of violence that, by virtue of its characteristics, is *inherently* systematic.

We argue that the process of killing with lethal autonomous weapon systems (LAWS) is always a systematized mode of violence in which all elements in the kill chain—from commander to operator to target—are subject to a technification. This technification incentivizes a moral devaluation of those targeted, while also degrading the moral agency of those involved in the application of autonomous violence. As a result, important restraints on the use of military force are jeopardized.

With this focus, the article builds on the extensive literature produced over the past decade critiquing the development and use of LAWS on moral grounds.⁵ This article advances this critique in two important ways. First, by situating LAWS within the *longue durée* of systematic killing, we more accurately draw out the similarities and dissimilarities between these systems and earlier modes of systematic violence. Such analysis is too often lacking from criticisms of this technology,

which typically either exaggerate or downplay both the material and moral novelty of this type of violence. Second, we add important nuance to the long-standing claim that LAWS dehumanize human targets. This dehumanization is real, we argue, but impacts the moral status of both the recipients and the dispensers of autonomous violence. In the case of dispensers, dehumanization operates alongside, and is compounded by, problematic effects of authorization and routinization. With LAWS, all three are amplified—a technification that erodes moral constraints.

These insights serve as a counterpoint to the recent proliferation of scholarship making the moral case in favor of LAWS as a potentially more humane, or otherwise “better,” alternative in the administration of lethal force.⁶ These supportive accounts, we argue, too often rest on an abstracted and overly idealistic concept of how the logic of LAWS operates within the broader setting of warfare. These perspectives neither take sufficient account of the wider historical dimension that underwrites the trajectory toward systematic killing nor adequately consider the real-world complexities and specificities of the technological system and its affordances in relation to killing in war. Correcting this oversight, we identify LAWS as both a continuation of and departure from the past, perpetuating historical processes of target and agent degradation while generating distinct and problematic technological specificities of systematic killing.

We begin the article with a brief overview of the AI-enabled lethal autonomous systems we are concerned with before tracing some of the key points in the current debates. We consider the motivations for acquiring these systems and the arguments put forward by proponents that they will ethically improve the battlefield. Section two engages the history of systematic killing, evaluating the degree to which such modes of violence influence and distort human relations and ethical considerations inside and outside the battlefield. In the final section, we explore how the factors that facilitate the erosion of moral restraint manifest in the processes prioritized by, and within, LAWS.

THE ALLURE OF AUTONOMOUS VIOLENCE

In this article, we are concerned with the relationship between systematization and violence in war. Before we engage with the types of lethal autonomous systems currently on the horizon, and the discourses associated with these systems and their human use, a very brief contouring of the concept of systematization is in order. Importantly, we do not argue that all systematic approaches to warfare are

problematic. Military organization and war fighting have been ordered and re-ordered throughout history into more fixed and instrumental systems.⁷ The rules of war have also been standardized to bind combatants to a more fixed set of proscriptive and prescriptive measures that limit the scope of permissible violence.

Our specific concern is with “intensified systematization”—modes of violence in which the logic of calculation, classification, and optimization for the act of elimination become paramount. This formulation of violence imperils essential moral restraints on the use of force and is intrinsic to AI-enabled lethal autonomous weapons. In this way, LAWS reproduce, and in some cases intensify, the moral challenges associated with prior episodes of intensified systematic killing.

Autonomous weapons technology has advanced significantly in recent years and is anticipated to continue doing so in the years ahead. Sophisticated AI innovations through neural networks and machine learning, paired with improvements in computer processing power, have opened up a field of possibilities for autonomous decision-making in a wide range of military applications, including the targeting of adversaries. Definitions of LAWS vary and remain hotly contested.⁸ The crucial aspect, however, is the weapon system’s potential to autonomously—without human intervention or action—select and engage targets. The definition offered by the International Committee of the Red Cross is widely used and offers a helpful delineation of autonomous weapons systems, and by extension LAWS: “Any weapon system with autonomy in its critical functions. That is, a weapon system that can select (search for, identify, track or select) and attack (use force against, neutralize, damage or destroy) targets without human intervention.”⁹ In contrast to remotely operated drones, LAWS relegate the human to a supervisory role in the kill chain loop (humans-on-the-loop), or remove the human entirely (humans-out-of-the-loop). In the latter case, targeting decisions and actions could be initiated and completed autonomously, based on input and sensor data, algorithms, and software programs.

Examples of LAWS include AI-enabled loitering munitions and AI-equipped, weaponized drone swarm systems that have the capacity to identify threats based on certain input parameters, fix on certain targets, and eliminate them once a threshold value has been reached. An AI-enabled weapon system like this would need to be trained on data that are relevant to a zone of conflict or area of engagement and require frequent updates, as “the introduction of new parameters or slightly heterogeneous data to the data under which the weapon has been trained will confound [LAWS].”¹⁰ The AI component of LAWS does

a significant amount of the independent cognitive work here. While several challenges arise from this highly dynamic process, for the commander as well as the operator, the allure of accelerated action in complex contexts with LAWS is strong. In theory—and in practice—LAWS can shorten the sensor-to-shooter timeline from minutes to seconds. The ability to navigate high complexity in an accelerated time frame is seen as a significant strategic benefit, even if it comes at the expense of direct human oversight. As General John Murray put it to a military academy audience in 2021, “Is it within a human’s ability to pick out which [swarm robots] have to be engaged?” and then make 100 individual decisions? “Is it even necessary to have a human in the loop?”¹¹ Statements like General Murray’s align with a broader vision of a fully networked, domain-crossing, and time-compressed AI-enabled future war.

Instructive here are the projects undertaken by the United States for a fully networked, domain-crossing “network of networks”—or a system of systems—that connects the data sensors and shooters of all U.S. military domains and allied militaries for greater speed and scale of operations, as articulated in the Joint-All-Domain Command and Control (JADC2) concept. The concept responds to the problem identified by some U.S. Department of Defense officials that “future conflicts may require decisions to be made within hours, minutes, or potentially seconds compared with the multiday process to analyze the operating environment and issue commands.”¹² It is a vision of war that is fully systematized in all its processes and operations, including lethal targeting, and in which both the speed and the scale envisioned clearly prioritize autonomous violence. AI-enabled LAWS will be instrumental in realizing these visions.

The debate over the legal, ethical, and political implications of autonomous weapons systems is protracted and ongoing. Seminal critical voices have urged a halt to the development and use of LAWS on account of their incommensurability with existing moral and legal standards in war.¹³ Noel Sharkey and Lucy Suchman, for example, argue that LAWS lack the technological sophistication and capabilities to adhere to the principle of distinction or proportionality, two core elements of international humanitarian law (IHL).¹⁴ Robert Sparrow and others have argued that AI-enabled weapons systems produce a responsibility gap in situations where the system makes an unexpected or unlawful lethal decision for which nobody can viably be held to account.¹⁵

More foundational, deontological objections to LAWS have also been voiced. Christof Heyns, for example, makes a strong case against autonomous violence

based on the fact that “humans should not be treated similar to an object that simply has instrumental value . . . or no value at all.” In the case of LAWS, Heyns writes, human targets have “no avenue, futile or not, of appealing to the humanity of the enemy, or hoping their humanity will play a role, because it is a machine on the other side.”¹⁶ Criticisms of this type frame autonomous violence as inherently immoral, on account of its violation of the principle of human dignity. Our discussion builds on these interventions, clarifying both the technological shortcomings of these systems and their problematic marginalization of human judgment and values.

Moral arguments also extend in the other direction, in favor of the development and use of LAWS. Deane Baker, for example, argues that LAWS ultimately reflect the intent of, and can be controlled by, those who decide to employ them. Thus LAWS can remain, in principle, “compliant with the ethics and laws of war.”¹⁷ In this line of argumentation, the human remains foregrounded as the only relevant moral actor. Baker objects to the idea that any kill decision is “ceded” or “delegated” to the machine and for him, and others, it is obvious that humans will always make the kill decision; LAWS are framed here as instruments—neutral tools to be used or misused discretionarily like any other weapons. Heller makes a similar point, noting that “[LAWS] do not ‘decide’ at all; they simply execute the targeting rules that humans have programmed into them.”¹⁸ These observations are true but incomplete, relying on an overly abstracted, highly idealized, and, in some cases, overly simplified version of autonomous weapons and the human in relation to such systems. As Thompson Chengeta has convincingly argued in response to such claims, “Where a machine is designed to make all critical decisions without human control, responsibility to make legal and ethical judgments has, in fact, been delegated to the machine.”¹⁹

The instrumental view of LAWS brackets the realities of machine-learning logics, including the fact that AI-enabled systems must be trained, not merely programmed; that a significant degree of unsupervised calculations are a key part of any AI-enabled system; that these systems rest on a logic of error and iteration, meaning improvements will often be paid for in lost lives; that frequent updates complicate any verification and validation process; and that it is highly unlikely that such systems can work as intended, consistently, within the messy complexities of any zone of conflict.²⁰

A related but different moral argument frames LAWS as a potentially lifesaving technology, as illustrated by the arguments put forward in the 2021 National

Security Commission on Artificial Intelligence's *Final Report*: "If properly designed, tested, and used," the argument goes, LAWS "could *improve* compliance with International Humanitarian Law."²¹ The idea here is that by using systems capable of capturing and processing larger amounts of data more accurately and at faster speeds, better decisions can be made, at the expense of fewer civilian casualties.²² As Arkin points out, as have others, "Unmanned robotic systems can be designed without emotions that cloud their judgment or result in anger and frustration with ongoing battlefield events."²³ Human cognition and emotion is framed as an impediment, not facilitator, of good conduct in war; LAWS are the corrective, enabling those empowered to "engage in unwavering, precise combat."²⁴ Cappuccio and colleagues take a similar line of reasoning in suggesting that LAWS "can relieve military personnel from the burden of killing, thus sparing them the risk of suffering moral injury, even if the available artificial intelligence (AI) is not sophisticated enough to solve complex ethical puzzles."²⁵

This is a compelling narrative, but it rests on a speculative and superficial understanding of the logical implications of this technology specifically, and systematic violence more generally. It also decontextualizes these weapons to a problematic degree. As Alexander Bellamy notes in his work on mass atrocities, "Arguments are not aired and received in a vacuum." Social and other contexts matter when it comes to understanding the dynamics of violence, including the "material and institutional power of the perpetrators."²⁶ LAWS proponents too often ground their optimism in overly abstract potentialities while ignoring the actual history of systematization in war, as well as the moral and legal records of those most likely to utilize this technology.

We ask, what if instead of preserving or improving upon the "goodness" of human military personnel, an intensified system logic facilitates a worsening of battlefield conduct? In the next section, we historicize this claim, detailing how intensified processes of systematic killing imperil the moral status of both the recipients and the dispensers of violence, to the detriment of essential restraints on military force.

SYSTEMATIC KILLING IN HISTORY

Much of the debate over the morality of autonomous weapons centers on the function and value of "humanity" in war. Opponents of these systems have been criticized for comparing autonomous weapons not to human combatants

as they most often are—confused, distressed, and volatile—but rather to human combatants as they wish them to be—reflective, rational, and compassionate. The “IHL compliant just warrior” image of humanity is an ideal type, at odds with much of the human experience in past and present war.²⁷ Human combatants internalize the rules of the battlefield too slowly and discard them too quickly for them to be consistently effective. For as long as war has been fought, human participants, driven by rage, fear, and hatred, have given in to their “mad passions” and terrorized and murdered innocent parties.²⁸ It is this image of humanity—as a cause, not a corrective, of misery in war—that proponents of autonomous weapons reference when framing the technology as an ethically superior alternative.²⁹

While not incorrect, this pessimistic depiction of humanity is excessively narrow, excluding other types of human-driven misconduct and immorality that autonomous weapons systems are more likely to accelerate than prevent. True, atrocity in war is often sourced in passion—hate of the enemy and exhilaration and joy in their suffering. However, alongside this are the more dispassionate and systematically dispensed cruelties. “Cold violence,” Jonathan Glover writes, “should disturb us far more than the beast of rage in man.”³⁰ These colder modes of killing are driven less by personal animus than a logic-driven calculation to extirpate. Systematic, dispassionate “pest control” models of killing have been a feature of some of the most destructive episodes of human history.³¹

Discomfort over systematic killing derives primarily from its historical association with inhumane and unjust harm; the harnessing of systems-oriented action in service of the mass killing of the undeserving. Importantly though, many are also repelled by the *process* of systematic killing and the degree to which subsuming human agency and emotion into intensified systems of violence undermines the moral status of both the dispensers and the recipients of harm. Drawing out this history and the empirical realities of violence in war helps us move beyond autonomous weapons debates overly infused with abstract theoretical assumptions.

Target Degradation

We argue that the processes associated with systematic killing, especially the more intensified versions, imperil restraints on the use of force. This can first be observed in relation to the status of those targeted. Systematization either directly imposes or incentivizes totalizing categories that suppress the individual differences of the targeted, including differences that might inform our moral judgment as to whether targeting is just.

These processes were at work in much of the colonial violence of previous centuries. As Lawrence Freedman notes, “Colonialism established the idea of whole populations as legitimate targets.”³² Categorization was fixed in this context, with targeted individuals denied the opportunity to express their innocence, and by extension, their immunity from direct and deliberate harm. British commander in chief Herbert Kitchener’s description of British tactics during the Second Boer War exemplifies this process. The British were to

flush out guerrillas in a series of systematic drives, organised like a sporting shoot, with success defined in a weekly “bag” of killed, captured and wounded, and to sweep the country bare of everything that could give sustenance to the guerrillas, including women and children. . . . It was the clearance of civilians—uprooting a whole nation—that would come to dominate the last phase of the war.³³

As the colonial case makes clear, the moral challenges implicit in the contemporary systematization of violence have a longer history.³⁴ Anchored as it is to fixed categorization, killing by a system logic greatly reduces, if not eliminates, the possibility of interpersonal connection, or even recognition. The objectified person on the receiving end of lethal force has little-to-no agency in the targeting process; no recourse to know how his or her data is dis- and re-aggregated in the production of the category “enemy object.” Within such systems, inferences are drawn and assumptions made that encase categories such as “enmity” in discrete terms.

This logic stems from, and feeds into, what Hannah Arendt described as the totalitarian ambition toward “knowing” the enemy based on data classification and cross-tabulation. Nazi policies were characterized by a process of objectification and dehumanization. The systematic classification of humans for elimination en masse severed the very premise for human relations—that of being considered as an individual: a subject, not an object. Under the SS corps, “Bestiality gave way to an absolutely cold and systematic destruction of human bodies; calculated to destroy human dignity,” and kill off any individuality of those imprisoned in the camps.³⁵ Inmates became objects, classified based on a system of identification “according to which each prisoner had a rectangular piece of material sewn onto his or her uniform” upon which a “classification triangle” was placed that indicated by color whether that person was categorized as a political prisoner, a Jehovah’s Witness, a prostitute or other “asocial,” a homosexual, a criminal, a Jew, and so on.³⁶ Harnessing new technologies to kill at a distance—both

physically and socially—and thus avoid the “horrors of face-to-face-killing,” the violence of Nazi Germany offers a stark example of the moral challenge of systematic, dispassionate violence.³⁷

Readers may understandably question the relevance of such examples for current debates over the dangers of autonomous weapons. We should firstly reiterate that we do not seek to draw moral equivalence between the genocidal practices of World War II and the use of distanced autonomous killing. What these historical examples do show, however, is systematic killing at its most pathological. Analysis of these cases allows us to better recognize the problematic features of systematization-as-process that operates elsewhere, albeit to far less severe degrees. In reality, the systematization of violence is a spectrum. It ranges from the routine and unproblematic to the murderous and genocidal. In between are a number of even more recent examples, including ones from armed conflict, of systematic killing that is nongenocidal but still radically in tension with prevailing moral standards.

One such example can be found in the systems-logics that governed U.S. conduct in Vietnam. In 1960s U.S. military doctrine—particularly under the direction of Robert McNamara, a Ford Motor Company executive turned secretary of defense—modes of warfare were forged along the lines of highly quantitative computational processes: James Gibson famously termed this “technowar.”³⁸ As Nick Turse writes:

[McNamara] relied on numbers to convey reality and, like a machine, processed whatever information he was given with exceptional speed, making instant choices and not worrying that such rapid-fire decision making might lead to grave mistakes. . . . McNamara and his national security technocrats were sure that, given enough data, warfare could be made completely rational, comprehensive, and controllable.³⁹

The implementation of this scientific computational management ethos translated into an undue focus on cost-benefit evaluations that assumed—rationally—that more deaths on the side of the enemy would spell victory for the United States. The relevant statistic was “body count,” which led to a mandate to kill as many suspected enemies as possible. Needless to say, quota-based killing is fraught with moral risk. In Vietnam, enemy classification was broadly and often crudely drawn: “Everyone in a conical hat or the loose-fitting Vietnamese clothes that Americans called ‘black pajamas’ was a potential adversary.”⁴⁰ This objectification facilitated and excused the commission of numerous acts of battlefield negligence

and atrocity by American forces during the conflict. The exact figures of combatant and noncombatant deaths on the side of the Vietnamese is unclear; despite the intense focus on data and body count, the innocent bodies were never fully counted. Various estimates suggest that there were between 1.1 million and 3.8 million violent deaths (civilians and combatants) and around 5.3 million wounded civilians.⁴¹ According to U.S. medic Wayne Smith, the “body-count” system led to “a real incentivizing of death and it just fucked up our value system.”⁴²

These same dangers endure in the algorithmic warfare of today. This danger is twofold. In addition to the target degradation, intensified systematization threatens the moral status of those who dispense violence.

Agent Degradation

A common question when examining systematic violence throughout history, particularly the more morally egregious examples, is how could they do it? How could individuals, not all of whom seem outwardly evil, contribute to a system of mass-produced murder? The answer to these questions can inform our understanding of the present and future dangers of autonomous killing.

The psychologist and scholar of atrocities Herbert C. Kelman offers one such answer. He recognizes that a “historically rooted and situationally induced” hostility—often along racialized lines—forms a substantive element in systematic mass killing, but Kelman argues that it is not a primary instigator for large-scale violence. Rather, he advises us to consider “the conditions under which the usual moral inhibitions against violence become weakened.” In his 1973 work on mass violence, he identifies “authorization,” “routinization,” and “dehumanization” as important contributors to this weakening of moral restraint.⁴³

“Authorization” provides the necessary substrate for sanctioned transgressions at scale. When a legitimate authoritative agent explicitly orders, implicitly encourages, or tacitly approves acts of violence, “people’s readiness to condone them is considerably enhanced.”⁴⁴ Through authorization, control is surrendered to authoritative agents bound to larger, often abstract goals that “transcend the rules of standard morality.”⁴⁵ For those tasked with the actual delivery of violence, agency is lost, or abdicated, to central authorities, who in turn cede their authority to still higher powers. This layered referral separates cognition from affect, and personal morality from a rationalized appeal to overriding violence.

The second process Kelman highlights in the erosion of moral restraints is “routinization.” Whereas authorization overrides otherwise existing moral concerns,

processes of routinization limit the points at which such moral concerns can, and will, emerge.⁴⁶ Routinization fulfills two functions: first, it reduces the necessity of decision-making, thus minimizing occasions in which moral questions may arise; and second, it makes it easier to avoid the implications of the action, since the actor focuses on the details rather than the meaning of the task at hand.⁴⁷

The third process, and the one that arguably connects most closely with the target objectification already discussed, is “dehumanization.” Processes of dehumanization work to deprive victims of their human status; “to the extent that the victims are dehumanized, principles of morality no longer apply to them and moral restraints against killing are more readily overcome.”⁴⁸ Importantly though, the same processes that degrade the moral status of the victim may also dehumanize perpetrators:

Through his unquestioning obedience to authority and through the routinization of his job, he is deprived of personal agency. He is not an independent actor making judgments and choices on the basis of his own values and assessment of the consequences. Rather, he allows himself to be buffeted about by external forces. He becomes alienated within his task.⁴⁹

This condition is pronounced within the digital logics of AI-enabled systems. Before we detail this, however, it is again important to reemphasize that problematic systematization is not specific to any one technology or mode of war. Alongside the examples already given, we can look to the U.S. armed drone program for a more recent illustration of the problematic effects of authorization, routinization, and dehumanization.

Within the U.S. drone program, armed drones were one part of a “flexible and persistent network of capabilities spanning global distance and woven together by arrays of streaming data.”⁵⁰ Within this system of integrated technologies, numerous moral challenges emerged, particularly in the context of targeted killing. Initially justified as a necessary response to “confirmed [terrorists] at the highest level,” targeting standards deteriorated as drone killing became more routinized.⁵¹ Over the course of the War on Terror, the United States, argues Ryan Devereaux, devoted “tremendous resources to kill[ing] off a never-ending stream of nobodies.”⁵² Obama himself made reference to the systematic nature of U.S. drone killing, and the moral slippage it incentivized:

The problem with the drone program . . . is that it starts giving you the illusion that it is not war . . . the machinery of it started becoming too easy, and I had to actually impose

internally a substantial set of reforms in the process to step back and remind everyone involved this isn't target practice.⁵³

Dehumanization was also a feature of the U.S. drone program, compounded by the data-driven nature of the killing, with those targeted sometimes likened to weeds and pests.⁵⁴ According to one American intelligence source, the internal view of the special operations community toward those hunted by armed drones was: "They have no rights. They have no dignity. They have no humanity to themselves. They're just a 'selector' to an analyst. You eventually get to a point in the target's life cycle that you are following them, you don't even refer to them by their actual name." This practice, he said, contributes to "dehumanizing the people before you've even encountered the moral question of 'is this a legitimate kill or not?'"⁵⁵

This section has highlighted a number of episodes of systematic killing, across a range of historical periods. They vary significantly, in terms of both the means and the ends of the violence in question. Within this variance, however, commonalities can be observed. Systematic violence, while not inherently problematic, generates inescapable moral challenges, particularly in cases of intensified systematization. This includes the erosion of moral status for both the dispensers and the recipients of violence. This loss has the potential to negatively impact restraint in war, a risk that endures today in the context of systematic autonomous killing.

THE MORAL CHALLENGE OF LAWS

New technologies can disrupt the status quo of war in different ways. In some cases, "disruptive" technological innovation does not create, but rather makes more salient, enduring but unresolved problems in war.⁵⁶ The moral challenge of autonomous violence is an example of this. Systematic killing in war precedes this technology and goes beyond it. Autonomous systems do, however, accelerate many of its worst features, by virtue of its particular technical characteristics.

Before expanding on these characteristics, it should be reiterated that the moral challenge of autonomous violence is not one of *inhumanity*. Humans will remain intrinsic to these systems—at issue is the *type* of humanity this technology makes less and more likely. Autonomous weapons, in delivering us from the passionate, volatile misconduct of human individuals, risk plunging us ever further into the cold, dispassionate misconduct of human systems.

Seeing like a Computer: The Human Object

Modes and systems of classification for the grouping and ordering of enemy categories are as old as warfare itself. We categorize and classify to give order to our actions and interactions with others, and we frequently use signifiers (such as uniform and insignia) to do so. However, as warfare has become more complex, geographically distributed, and asymmetrical, traditional identifiers no longer render the enemy coherently legible and visible, and increasingly data serves as a stand-in. As Josef Ansorge writes, “Under such challenging conditions of illegibility and disfluency . . . data is sought to an unprecedented degree” to identify and track enemies and predict who might become one.⁵⁷

This logic is amplified with LAWS, where AI systems understand and identify targets based purely on object recognition and classification via neural networks. AI renders the world as it perceives the world, as a set of objects and related patterns from which outcomes can be predicted and calculated, including the decision over which “objects” are to be targeted. Why an individual might be marked for elimination might have little to do with who they are, how they behave, or what they intend. Rather, the target comes to be known through statistical probability, wherein “seemingly discrete, unconnected phenomena are conjoined and correlatively evaluated.”⁵⁸ Within this process, data—behavioral, contextual, image, perhaps medical, and so on—are disaggregated and reaggregated to conform to specific modes of classification. Drawing upon this data, the system calculates a systematic inference of who, or what, falls within a pattern of normalcy (benign) or abnormality (potential threat) in order to eliminate the threat.

This form of enemy identification is fraught with the risk of seeing patterns and drawing inferences where there are none—a well-known challenge in human reasoning that becomes “baked” into algorithmic structures and systematized. An AI system tasked with image recognition, for example, “understands” an image as a set of pixels, and each pixel as a set of fields—that is, an “array of numbers, corresponding to the brightness and color of the image’s pixels.”⁵⁹ In order to train such systems for the purpose of identifying an enemy, the system would first need to be trained on a sizable number of appropriately labeled images (for example, for the category “terrorist” or “enemy”) fed into it. Through convolutional networks, certain image features are established as “useful for classifying the object it is trained on.”⁶⁰ This information is then fed into a neural network, which assorts and classifies the input to predict what *object* the image depicts with a

certain degree of confidence, expressed in percentage values. This is always an approximation, never a true and complete reflection of reality.

As Paddy Walker notes, “It is systemically difficult for LAWS to classify an event or object into a particular category. Its processes will instead review and dissect it according to an inappropriately small number of characteristics.”⁶¹ To remain relevant for the dynamic context of warfare, the weapon must “continually [calculate] new probabilities for its immediate world,” a process that is “governed by an error function.”⁶² In other words, a kill decision with LAWS is one that rests on approximation, streamlining, and a smoothing out of data points. Within this process, persons become not just objects in the selective application of violence, but objects that are constituted through algorithmic patterns. Patterns are identified and lines of association are drawn (where there possibly are none) and based on this, kill calculations are made. The very logic of AI rests on this classification and codification of life into computable data to identify objects, and patterns between objects. As John Cheney-Lippold notes, “To be intelligible to a statistical model is . . . to be transcoded into a framework of objectification” and become defined, cross-calculated, as a computationally ascertained, actionable object.⁶³ This epistemic grounding produces not only a pure objectification but also, if the target is human, a desubjectification and deindividualization. Such individuals “cannot rely on anything unique to them because the solidity of their subjectivity is determined wholly outside of one’s self, and according to whatever gets included within a reference class or dataset they cannot decide.”⁶⁴

These logics of objectification rest at the heart of most human dignity–centric critiques of autonomous killing. The concept of “human dignity” is an ambiguous one, particularly when considered in the context of armed conflict. Designating *particular* weapons as “inhumane” may strike some as counterintuitive, given that all weapons, by design, injure and kill humans, often in painful and gruesome ways. Does it really make sense to label an autonomous weapon morally worse than a remotely operated drone if both carry identical payloads that create the same material effect: burnt, torn, and destroyed bodies? Focusing on the material effect of the strike, however, misses the moral stakes at issue. The human dignity challenge of LAWS stems not from the material character of the targeting *effect*, but rather from the moral character of the targeting *process*. Humans deserve to remain free of predation from systems that lack the capacity to properly assess and weigh their moral worth. LAWS fail to meet this standard by virtue of their inability to recognize and act upon the full range of factors that render an

individual morally liable—or not—to lethal targeting. This is especially true when LAWS are proposed for counterterrorism and urban warfare scenarios, where they would be necessarily tasked with significantly more complex duties than the object recognition of uniformed adversaries.

Computational objectification of the sort LAWS produce necessarily dehumanizes the targeted individual. It reclassifies humans as *something* less; *something* that is statistically determined, processed, and rendered actionable. LAWS have no conception, by design, of a subject-object relation, as we humans would, nor do they understand the subjective self in relation to the object upon which they act. It is in *this* absence, not in the Hellfire missiles (or whatever else) used by the platform to end the life of the targeted, where the challenge to human dignity resides.

The Technification of Being

It is not only the targeted who are negatively impacted by systematic killing. Participant agency may also be reduced, or overwhelmed, by the systems logic that governs the organization and infliction of violence. Within a context of systematic killing, many of the features we would wish to cultivate and preserve in war—judgment, personal responsibility, self-reflection, moral restraint, and so on—have little space to operate. What enters instead is a logic of efficiency and speed in which the human—whether that is the commander or the operator—is tasked to work within the respective system logic. This challenge is evident in autonomous violence, a method of killing that generates, like its historical antecedents, problems related to authorization, routinization, and dehumanization.

Authorization

As highlighted in the second section, Kelman refers to structures of authority as one of the prongs that license and sustain mass killing. The computational systems of LAWS fit logically into this scheme of abstraction and abdicated responsibility. Computational systems command deference from operators or commanders who rarely fully understand the processes involved in the computational decision. Within conditions of such complexity and abstraction, humans are left with little choice but to trust in the cognitive and rational superiority of this clinical authority. This relationship is often comprised under the term “automation bias” and is a well-documented phenomenon in the literature on human-machine interactions,⁶⁵ whereby technological authority serves to smooth over moral tensions.⁶⁶ At issue with LAWS is not a formal, hierarchical process of authorization, but

rather one that emerges from the ostensibly neutral and superior character of the machine itself. Regardless of rank, the ability—and possibly willingness—to challenge the authority of the machine logic becomes weakened.

When placed within a complex digital environment, human cognition, experience, and action are mediated and moderated through machine logics. Within such a framework, the possibility to exercise moral agency is significantly truncated for both commanders and operators.⁶⁷ In other words, agency is affected across the distributed control setting relevant to human control of LAWS.⁶⁸ The effect is, however, particularly pronounced when operators are called upon to action incoming information—including information of life and death stakes—within seconds. Operators burdened by such constraints may lack both a “sufficient level of situational awareness to make meaningful judgements” and “sufficient insights into the parameters under which the automated or autonomous parts of the command modules select and prioritise threats to scrutinise target selection and . . . abort the attack.”⁶⁹ As operators on the loop, the human thus becomes the .exe module⁷⁰ in the wider computational network, with only limited—if any—capacity to override or intervene into the preset action. The combination of a commander deciding to deploy LAWS and the perceived superiority of the machine logic is highly likely to yield a context in which military personnel become “involved in an action without considering the implications of that action and without really making a decision.”⁷¹

This is not to argue that humans are a foolproof safeguard against wrongdoing in war; they very clearly are not. But critically, this does not rid LAWS of their morally problematic aspects. We should prefer conditions where those charged with doing violence understand the context and consequences of their actions, are able to recognize when they should relent from violence, and have the ability to act upon this impulse rather than becoming removed from the process. The danger of LAWS is not that they will too often fall short of these standards—more human-centric systems do this constantly—it is that they will lack the very *capacity* to meet these standards.

Routinization

Routinization operates at both the individual and the organizational level, shifting the focus onto the purely procedural. At the individual level, the operator, as a functional element in the system’s logic, focuses on the specific executive task at hand with limited situational overview. At the organizational level, tasks relating

to the action (system design, algorithm programming, setting parameters for system action, and so forth) are divided and often diffuse. The process of routinization facilitates efficiency, procedural accuracy, speed in executing the task, and so on. This, in turn, becomes the norm, the standard for conducting the action well, and “the nature of the task becomes completely dissociated from their performance of it.”⁷² The primary danger of the routinization of violence is that it will foreclose opportunities for moral intervention and thus weaken moral restraint.⁷³

In conflict, there are always ambiguities that remain unresolved, where certainty cannot be established as to the identity and liability of a potential target. Even when we possess a set of parameters to make such determinations with reasonable confidence, some degree of uncertainty endures. It is precisely these ambiguities that leave space for ethical reasoning, which, in turn, allows for ethical intervention. As we have seen above, such interventions are necessary when the system, or indeed the rules, are either structurally or episodically overinclusive, mandating the targeting of those who have been categorized falsely as legitimate targets. LAWS, framed by some as having the potential for “ethical prowess,”⁷⁴ divorce cognition from emotion, leaving us with less morally empowered agents of violence.

Importantly, this challenge is not limited to the violent end of the kill chain. The computerized routinization built into LAWS narrows the space for human agency. Humans remain within the system, but responsibility for lethal force—the parameters and formulation of objections, as well as the execution of violence—is diffused, or detached, through the systems process. At its worst, this detachment of responsibility facilitates the careless or indeed deliberate application of wrongful violence.

Dehumanization

As indicated above, where a systematic approach to killing is applied, dehumanization is often twofold. Targets are objectified and stripped of the rights and recognition they would otherwise be owed by virtue of their status *as* humans. They are then reinterpreted as something less, “*something* that needs killing.”⁷⁵ This process also, however, typically dehumanizes the perpetrator. The dehumanization of the soldier, the operator, and those that set the parameters for killing takes hold gradually as he/she functions within the wider system of killing in which cognition and affect become starkly separated. Where personal responsibility, human

relations, and empathy are systematically discarded, one cannot act as a human moral being.⁷⁶ As the rich literature on this topic reveals, these processes invite violence and abuse.⁷⁷

We may again ask whether these challenges are inherent to the system or (potentially) resolvable. According to some, many of the problems we associate with LAWS can be mitigated if we make humans functionally relevant within the system; more fully lean into the *technification of being*, in other words. Humans, the argument goes, must be able to understand and work with the logic of the machine for a superior outcome: “*Weak human + machine + better process [is] superior to [either] a strong computer alone [or] . . . a strong human + machine + inferior process.*”⁷⁸ This is intuitive for future warfare only if we accept that the system logic should prevail in the process of killing with LAWS. Such a future would give sanction to systems of violence that cast enemies, inescapably, as inhuman objects, and render combatants ever more morally inert. LAWS accelerate and rationalize the decision to kill, but, in doing so, open up new spaces for moral infraction. Those concerned with the regulation and “humanization” of war should look elsewhere than the computational and dispassionate violence of autonomous weapons.

CONCLUSION

“Nothing is so dangerous in the study of war,” argued Sir Julian Corbett, “as to permit maxims to become a substitute for judgment.”⁷⁹ This warning is equally applicable to processes of systemization, which subsume human judgment on the battlefield to a morally problematic degree. As highlighted above, international humanitarian law is valued precisely because human judgment has proven throughout history to be an insufficient check on individual conduct. The principle of discrimination, the prohibition against perfidy, and prisoner of war protections—these rules stand whatever situational pressures exist and regardless of the vagaries of combatant judgment. Our appreciation for these rules should not, however, blind us to the dangers of the inverse—cold and dispassionate forms of systematic violence that erode the moral status of human targets, as well as the status of those who participate within the system itself.

The argument that LAWS can be more ethical agents in war can only hold if we think of war as a largely procedural and process-focused activity in which moral lines are relatively easy to identify and sufficiently robust to withstand uncertainty

and ambiguity. This might be an ideal, but it is not, and likely never will be, a reality. War is riven by a complexity that precludes certainty; and by extension, the smooth and reliable application of systematic violence to target objects. To proceed as if this is not the reality, to impose systematic violence upon environments structurally unsuited to such an approach, is to court foreseeable and ruinous moral harm.

With LAWS, and AI-infused killing more broadly, violence becomes systematized in the most literal sense. The system provides the organization, optimized function, distancing, and moral vacuum required to expand modes of killing rather than fostering restraint. This is neither genocide nor ethnic cleansing nor any of the other forms of historical systematic murder examined in this article. The violence of LAWS is not morally close to the mass killing that punctuated so much of the twentieth century. What we do observe, however, is an echo of the problematic past in the autonomous *processes* of today: an implicit set of conditions that might facilitate moral infraction in the use of lethal violence in warfare.

NOTES

¹ Authors listed alphabetically.

² *Burning Conscience: The Case of the Hiroshima Pilot Claude Eatherly, Told in His Letters to Günther Anders*, letter from Günther Anders to Claude Eatherly (London: Weidenfeld & Nicolson, 1962), p. 1.

³ *Ibid.*, pp. 108–9.

⁴ Norbert Wiener, *The Human Use of Human Beings: Cybernetics and Society* (London: Free Association Books, 1989) (emphasis in the original), p. 185.

⁵ Christof Heyns, “Autonomous Weapons in Armed Conflict and the Right to a Dignified Life: An African Perspective,” *South African Journal of Human Rights* 33, no. 1 (April 2017), pp. 46–71; and Peter Asaro, “On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making,” *International Review of the Red Cross* 94, no. 886 (June 2012), pp. 687–709.

⁶ See, for example, Steven Umbrello, Phil Torres, and Angelo F. De Bellis, “The Future of War: Could Lethal Autonomous Weapons Make Conflict More Ethical?,” *AI & Society* 35 (March 2020), pp. 273–82; Erich Riesen, “The Moral Case for the Development and Use of Autonomous Weapon Systems,” *Journal of Military Ethics* 21, no. 2 (2022), pp. 132–50; Garry Young, “Should Autonomous Weapons Need a Reason to Kill?,” in “Honour and Admiration after War and Conflict,” special issue, *Journal of Applied Philosophy* 39, no. 5 (November 2022), pp. 886–900; Duncan MacIntosh, “Fire and Forget: A Moral Defense of the Use of Autonomous Weapons Systems in War and Peace,” in Jai Galliott, Duncan MacIntosh, and Jens David Ohlin, eds., *Lethal Autonomous Weapons: Re-Examining the Law and Ethics of Robotic Warfare* (New York: Oxford University Press, 2021), pp. 9–24; Massimiliano Lorenzo Cappuccio, Jai Christian Galliott, and Fady Shibata Alnajjar, “A Taste of Armageddon: A Virtue Ethics Perspective on Autonomous Weapons and Moral Injury,” *Journal of Military Ethics* 21, no. 1 (2022), pp. 19–38; Kevin Jon Heller, “The Concept of ‘The Human’ in the Critique of Autonomous Weapons,” *Harvard National Security Journal* (forthcoming 2023); and Ryan Jenkins and Duncan Purves, “Robots and Respect: A Response to Robert Sparrow,” *Ethics & International Affairs* 30, no. 3 (Fall 2016), pp. 391–400.

⁷ John Keegan, *A History of Warfare* (London: Pimlico, 1994).

⁸ A lack of definitional coherence has been frequently cited as a regulatory roadblock throughout the international debate over the legality of LAWS. See Mariarosaria Taddeo and Alexander Blanchard,

- “A Comparative Analysis of the Definitions of Autonomous Weapons,” *Science and Engineering Ethics* 28, no. 37 (August 2022), pp. 1–22.
- ⁹ International Committee of the Red Cross, “Views of the International Committee of the Red Cross (ICRC) on Autonomous Weapon Systems” (Convention on Certain Conventional Weapons [CCW], Meeting of Experts on Lethal Autonomous Weapons Systems [LAWS], Geneva, April 11–15, 2016), p. 1, www.icrc.org/en/download/file/21606/ccw-autonomous-weapons-icrc-april-2016.pdf.
- ¹⁰ Paddy Walker, “Leadership Challenges from the Deployment of Lethal Autonomous Weapon Systems,” *RUSI Journal* 166, no. 1 (May 2021), pp. 10–21, at p. 14.
- ¹¹ John Murray, quoted in Will Knight, “The Pentagon Inches toward Letting AI Control Weapons,” *WIRED*, May 10, 2021, www.wired.com/story/pentagon-inches-toward-letting-ai-control-weapons/.
- ¹² Congressional Research Service, “Defense Primer: Command and Control,” “In Focus,” updated November 14, 2022, p. 2.
- ¹³ Asaro, “On Banning Autonomous Weapon Systems,” p. 688.
- ¹⁴ Noel Sharkey and Lucy Suchman, “Wishful Mnemonics and Autonomous Killing Machines,” *AISB Quarterly* 136 (May 2013), pp. 14–22.
- ¹⁵ Robert Sparrow, “Killer Robots,” *Journal of Applied Philosophy* 24, no. 1 (2007), pp. 62–77. For a comprehensive overview of the various perspectives in the debate on the responsibility gap, see Ann-Katrien Oimann, “The Responsibility Gap and LAWS: A Critical Mapping of the Debate,” *Philosophy & Technology* 36, no. 1 (January 2023), pp. 1–22.
- ¹⁶ Heyns, “Autonomous Weapons in Armed Conflict and the Right to a Dignified Life,” p. 63. See also Elvira Rosert and Frank Sauer, “Prohibiting Autonomous Weapons: Put Human Dignity First,” *Global Policy* 10, no. 3 (September 2019), pp. 370–75; and Daniele Amoroso and Guglielmo Tamburrini, “The Ethical and Legal Case against Autonomy in Weapons Systems,” *Global Jurist* 18, no. 1 (September 2017), p. 63 (for both quotes), www.degruyter.com/document/doi/10.1515/gj-2017-0012/html?lang=en.
- ¹⁷ Deane Baker, *Should We Ban Killer Robots?* (Cambridge, U.K.: Polity, 2022), p. 111; and Heller, “The Concept of ‘the Human’ in the Critique of Autonomous Weapons.”
- ¹⁸ Heller, “The Concept of ‘the Human’ in the Critique of Autonomous Weapons,” p. 5.
- ¹⁹ Thompson Chengeta, “A Critique of the Canberra Guiding Principles on Lethal Autonomous Weapon Systems,” E-International Relations, April 15, 2020, www.e-ir.info/2020/04/15/a-critique-of-the-canberra-guiding-principles-on-lethal-autonomous-weapon-systems.
- ²⁰ Walker, “Leadership Challenges from the Deployment of Lethal Autonomous Weapon Systems.”
- ²¹ National Security Commission on Artificial Intelligence, *Final Report* (Washington, D.C.: National Security Commission on Artificial Intelligence, March 2021), p. 91 (italics added), www.nscai.gov/2021-final-report.
- ²² Close to half of all U.S.-caused civilian casualties in the war in Afghanistan, for example, were the result of human misidentification. See Larry Lewis, *Redefining Human Control: Lessons from the Battlefield for Autonomous Weapons* (Arlington, Va.: Center for Autonomy and AI, Center for Naval Analysis, 2018).
- ²³ Ronald Arkin, “Lethal Autonomous Systems and the Plight of the Non-Combatant,” *AISB Quarterly* 137 (July 2013), p. 3.
- ²⁴ Umbrello et al., “The Future of War,” p. 276.
- ²⁵ Cappuccio et al., “A Taste of Armageddon,” p. 20.
- ²⁶ Alexander J. Bellamy, *Massacres and Morality: Mass Atrocities in an Age of Civilian Immunity* (Oxford: Oxford University Press, 2012), p. 384.
- ²⁷ John Williams, “Locating LAWS: Lethal Autonomous Weapons, Epistemic Space, and ‘Meaningful Human’ Control,” *Journal of Global Security Studies* 6, no. 4 (December 2021), pp. 1–18, at p. 9.
- ²⁸ Geoffrey Best, *Humanity in Warfare: The Modern History of the International Law of Armed Conflicts* (London: Methuen, 1980).
- ²⁹ Riesen, “The Moral Case for the Development and Use of Autonomous Weapon Systems,” p. 133.
- ³⁰ Jonathan Glover, *Humanity: A Moral History of the Twentieth Century* (London: Pimlico, 2001), p. 64.
- ³¹ Herfried Münkler, *Die neuen Kriege* (Hamburg: Rowohlt, 2003), p. 30.
- ³² Lawrence Freedman, *The Future of War: A History* (St. Ives, U.K.: Allen Lane, 2017), p. 36.
- ³³ Herbert Kitchener, quoted in Thomas Pakenham, *The Boer War* (London: Weidenfeld & Nicolson, 1979), p. 493.
- ³⁴ We examine this contemporary systematization in section three.
- ³⁵ Hannah Arendt, *The Origins of Totalitarianism* (New York: Random House, 2004), p. 585.
- ³⁶ Robert Jay Lifton, *The Nazi Doctors: Medical Killing and the Psychology of Genocide* (New York: Basic Books, 2000), p. 153.
- ³⁷ *Ibid.*, p. 162.

- ³⁸ James W. Gibson, *The Perfect War: Technowar in Vietnam* (New York: Atlantic Monthly Press, 1986). See also Nick Turse, *Kill Anything That Moves: The Real American War in Vietnam* (New York: Metropolitan Books, 2013), p. 38; and Antoine Bousquet, *The Scientific Way of Warfare: Order and Chaos on the Battlefields of Modernity* (New York: Columbia University Press, 2009), p. 149.
- ³⁹ Turse, *Kill Anything That Moves*, p. 38.
- ⁴⁰ *Ibid.*, p. 27.
- ⁴¹ *Ibid.*, p. 16–17.
- ⁴² Wayne Smith, quoted in *ibid.*, p. 45.
- ⁴³ Herbert C. Kelman, “Violence without Moral Restraint: Reflections on the Dehumanization of Victims and Victimizers,” *Journal of Social Issues* 29, no. 4 (Fall 1973), pp. 25–61, at p. 52.
- ⁴⁴ *Ibid.*, p. 39.
- ⁴⁵ *Ibid.*, p. 44.
- ⁴⁶ *Ibid.*, p. 48.
- ⁴⁷ *Ibid.*, p. 46.
- ⁴⁸ *Ibid.*, p. 48.
- ⁴⁹ *Ibid.*, p. 51.
- ⁵⁰ Timothy P. Schultz, “Remote Warfare: A New Architecture of Air Power,” in Phil M. Haun, Colin F. Jackson, and Timothy P. Schultz, eds., *Air Power in the Age of Primacy: Air Warfare since the Cold War* (Cambridge, U.K.: Cambridge University Press, 2022), pp. 26–53, at p. 28.
- ⁵¹ John Kerry, quoted in “U.S. Drone Programme: ‘Strict, Fair and Accountable’—Kerry,” News, BBC, May 28, 2013, www.bbc.com/news/av/world-radio-and-tv-22690918/us-droneprogramme-strict-fair-and-accountable-kerry.
- ⁵² Ryan Devereaux, “Manhunting in the Hindu Kush: Leaked Documents Detailing a Multi-Year U.S. Military Campaign in Afghanistan Reveal the Strategic Limits and Startling Human Costs of Drone Warfare,” *Intercept*, October 15, 2015, theintercept.com/drone-papers/manhunting-in-the-hindu-kush.
- ⁵³ Barack Obama, “Barack Obama: The 2020 60 Minutes Interview,” interview by Scott Pelley, *CBS News*, November 16, 2020, www.cbsnews.com/video/barack-obama-60-minutes-2020-11-15.
- ⁵⁴ Ed Pilkington, “Life as a Drone Operator: ‘Ever Step on Ants and Never Give It Another Thought?’,” *Guardian*, November 19, 2015, www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada.
- ⁵⁵ Quoted in Jeremy Seahill, “The Assassination Complex,” *Intercept*, October 15, 2015, theintercept.com/drone-papers/the-assassination-complex. This dehumanization functioned primarily at the policy, rather than the combat, level. There also exist numerous accounts of drone use that emphasize the intimacy and psychological strain of such strikes. See, for example, John Williams, “Distant Intimacy: Space, Drones, and Just War,” *Ethics & International Affairs* 29, no. 1 (Spring 2015), pp. 93–110.
- ⁵⁶ For a detailed overview of these distinctions in a legal context, see Rebecca Crootof, “Regulating New Weapons Technology,” in Eric Talbot Jensen and Ronald T. P. Alcalá, eds., *The Impact of Emerging Technologies on the Law of Armed Conflict* (New York: Oxford University Press, 2019), pp. 3–26.
- ⁵⁷ Josef Teboho Ansoorge, *Identify & Sort: How Digital Power Changed World Politics* (London: C. Hurst & Co., 2016), p. 124.
- ⁵⁸ John Cheney-Lippold, “Accidents Happen,” *Social Research: An International Quarterly* 86, no. 2 (Summer 2019), pp. 513–35, at p. 523.
- ⁵⁹ Melanie Mitchell, *Artificial Intelligence: A Guide for Thinking Humans* (London: Penguin Books, 2019), p. 78.
- ⁶⁰ *Ibid.*, pp. 86–87.
- ⁶¹ Walker, “Leadership Challenges from the Deployment of Lethal Autonomous Weapon Systems,” p. 17.
- ⁶² *Ibid.*, pp. 17, 14.
- ⁶³ Cheney-Lippold, “Accidents Happen.”
- ⁶⁴ *Ibid.*, pp. 523–24.
- ⁶⁵ See, for example, M. L. Cummings, “Automation Bias in Intelligent Time Critical Decision Support Systems” (AIAA [American Institute of Aeronautics and Astronautics] 3rd Intelligent Systems Conference, Chicago, September 20–22, 2004); and Noel Sharkey, “Guidelines for the Human Control of Weapons Systems,” (working paper for the CCW GGE, International Committee for Robot Arms Control, April 2018).
- ⁶⁶ Neta C. Crawford, “Bugsplat: US Standing Rules of Engagement, International Humanitarian Law, Military Necessity, and Noncombatant Immunity,” in Anthony F. Lang Jr., Cian O’Driscoll, and John Williams, eds., *Just War: Authority, Tradition, and Practice* (Washington, D.C.: Georgetown University Press, 2013), pp. 231–50.
- ⁶⁷ Elke Schwarz, “Autonomous Weapon Systems, Artificial Intelligence and the Problem of Meaningful Human Control,” *Philosophical Journal of Conflict and Violence* V, no. 1 (2021), pp. 53–72.

- ⁶⁸ For a discussion of distributed control related to the concept of meaningful human control of LAWS, see Merel Ekelhof, "Moving beyond Semantics on Autonomous Weapons: Meaningful Human Control in Operation," *Global Policy* 10, no. 3 (September 2019), pp. 343–48.
- ⁶⁹ Ingvild Bode and Tom Watts, *Meaning-Less Human Control: Lessons from Air Defence Systems on Meaningful Human Control for the Debate on AWS* (Odense: Center for War Studies, University of Southern Denmark, February 2021), p. 28.
- ⁷⁰ In programming, an .exe file is known as an executable file. Such a file contains coded instructions for a computational process that are activated by a user or another event.
- ⁷¹ Kelman, "Violence without Moral Restraint," p. 46.
- ⁷² *Ibid.*, p. 47.
- ⁷³ John Emery discusses these opportunities for moral engagement within any system of war as "meaningful inefficiencies," which makes space for important ethical deliberation. See John R. Emery, "Algorithms, AI, and Ethics of War," *Peace Review* 33, no. 2 (2022), pp. 205–12.
- ⁷⁴ Umbrello et al., "The Future of War," p. 276.
- ⁷⁵ David Livingstone Smith, *On Inhumanity: Dehumanization and How to Resist It* (New York: Oxford University Press, 2020), p. 9.
- ⁷⁶ Kelman, "Violence without Moral Restraint," p. 52.
- ⁷⁷ See, for example, Linda LeMoncheck, *Dehumanizing Women: Treating Persons as Sex Objects* (Totowa, N.J.: Rowman & Allanheld, 1985); Martha C. Nussbaum, "Objectification," *Philosophy and Public Affairs* 24, no. 4 (Autumn 1995), pp. 249–91; and Eileen L. Zurbruggen, "Objectification, Self-Objectification, and Societal Change," *Journal of Social and Political Psychology* 1, no. 1 (December 2013), pp. 188–215.
- ⁷⁸ Gary Kasparov, quoted in Trevor Phillips-Levine, Michael Kanaan, Dylan "Joose" Phillips-Levine, Walker D. Mills, and Noah Spataro, "Weak Human, Strong Force: Applying Advanced Chess to Military AI," War on the Rocks, July 7, 2022, warontherocks.com/2022/07/weak-human-strong-force-applying-advanced-chess-to-military-ai.
- ⁷⁹ Julian S. Corbett, *Some Principles of Maritime Strategy* (London: Longmans, Green & Co., 1911), p. 167.

Abstract: Systematic killing has long been associated with some of the darkest episodes in human history. Increasingly, however, it is framed as a desirable outcome in war, particularly in the context of military AI and lethal autonomy. Autonomous weapons systems, defenders argue, will surpass humans not only militarily but also morally, enabling a more precise and dispassionate mode of violence, free of the emotion and uncertainty that too often weaken compliance with the rules and standards of war. We contest this framing. Drawing on the history of systematic killing, we argue that lethal autonomous weapons systems reproduce, and in some cases intensify, the moral challenges of the past. Autonomous violence incentivizes a moral devaluation of those targeted and erodes the moral agency of those who kill. Both outcomes imperil essential restraints on the use of military force.

Keywords: war, armed conflict, ethics, technology, violence, autonomy, military AI, LAWS, lethal autonomous weapons systems, dehumanization