Human Dignity and Machine Agency: Reconciling International Humanitarian Law with the Ontology of Artificial Decision-Making

Sergey Sayapin[*]

ABSTRACT

As artificial intelligence (AI) increasingly mediates battlefield decision-making, international humanitarian law (IHL) faces a structural challenge – preserving its human-centred foundations in an algorithmic operational environment. This article introduces dignitary design as a jurisprudential framework for embedding respect for human dignity within the architecture, governance, and deployment of military AI systems. Grounded in the principle of humanity and the Martens Clause, dignitary design reconceptualises dignity as both a legal standard and a design requirement. The article argues that the legitimacy of future armed conflict regulation depends on integrating humanitarian values directly into technological systems, rather than relying solely on post hoc accountability.

Keywords: accountability – artificial intelligence (AI) – dignitary design – human dignity – international humanitarian law (IHL) – Martens Clause – military AI systems

I. Introduction: From Human Command to Algorithmic Mediation

Contemporary armed conflict is undergoing a profound structural transformation.[1] Artificial intelligence (AI) systems, once confined to auxiliary functions such as logistics optimisation, battlefield sensing, and intelligence analysis, are increasingly being integrated into operational processes that shape targeting, surveillance, force allocation, and strategic planning.[2] This development has generated growing concern among States, international organisations, military practitioners, and scholars regarding the future relationship between emerging technologies and the law of armed conflict.[3] As algorithmic systems become more deeply embedded in military decision-making, international humanitarian law (IHL) confronts a

[*]Dr Sergey Sayapin [s.sayapin@kimep.kz] is Professor of Law at KIMEP University (Almaty, Kazakhstan) and Distinguished Visiting Global Scholar at the NUS Centre for International Law (2025). The author thanks Dr. Samuel White and Dr. Jonathan Kwik for their helpful comments on an earlier draft of this article.

[1] ICRC, "International Humanitarian Law and the Challenges of Contemporary Armed Conflicts", *International Review of the Red Cross*, Vol. 106, No. 927, 2024.

[2] Wen Zhou and Anna Rosalie Greipl, "Artificial Intelligence in Military Decision-Making: Supporting Humans, not Replacing Them", *Humanitarian Law & Policy*, 29 August 2024, available at: https://blogs.icrc.org/law-and-policy/2024/08/29/artificial-intelligence-in-mil itary-decision-making-supporting-humans-not-replacing-them/ (all internet references accessed on 9 July 2026).

[3] ICRC and Geneva Academy, *Expert Consultation Report on AI and Related Technologies in Military Decision-Making on the Use of Force in Armed Conflicts*, ICRC, Geneva, March 2024.

fundamental challenge: how can a legal regime built upon assumptions of human judgment, responsibility, and moral agency preserve its normative integrity when increasingly consequential decisions are mediated by computational systems?

Debates concerning military AI have expanded rapidly over the past decade. Much of the existing literature has focused on the regulation of lethal autonomous weapons systems (LAWS), the requirement of meaningful human control, [4] accountability for algorithmically mediated harms,[5] and the adequacy of existing weapons-review mechanisms under Article 36 of Additional Protocol I.[6] These discussions have made important contributions to understanding the legal and ethical implications of emerging technologies. Yet they have tended to focus on questions of compliance, attribution, and regulation, often treating AI as an external technological development to which existing legal frameworks must adapt. Comparatively less attention has been devoted to a more fundamental question: what happens to the human-centred normative structure of IHL when operational judgment itself becomes increasingly mediated by algorithmic systems?

Addressing this question requires greater conceptual precision regarding the technologies under discussion. Not every form of automation raises the concerns examined in this article. Military operations have long relied on technologies that mediate perception, communication, navigation, targeting, and damage assessment. Automated defensive systems, precision-guided munitions, and collateral damage estimation tools have been integral components of military decision-making for decades. The challenge posed by contemporary AI therefore does not arise from technological mediation as such – rather, it emerges when algorithmic systems begin to perform evaluative, predictive, and classificatory functions traditionally associated with human judgment.[7] Machine-learning systems capable of identifying patterns, generating recommendations, prioritising targets, predicting civilian presence, or shaping proportionality assessments do not merely execute pre-programmed instructions.[8] Instead, they increasingly participate in the production of operational knowledge upon which legally significant decisions are based. Thus, the legal significance of AI lies not in automation alone but in its growing role in structuring the informational and evaluative environment within which decisions concerning the use of force are made.

---

[4] Daniele Amoroso and Guglielmo Tamburrini, "Autonomous Weapons Systems and Meaningful Human Control: Ethical and Legal Issues", *Current Robotics Reports*, Vol. 1, 2020.
[5] Gerald Mako, "Legal Accountability for AI-Driven Autonomous Weapons", *Articles of War*, 9 March 2026, available at: https://lieber.westpoint.edu/legal-accountability-ai-driven-autonomous-weapons/.
[6] Natalia Jeglevskaja, *International Law and Weapons Review: Emerging Military Technology under the Law of Armed Conflict*, Cambridge University Press, Cambridge, 2021, pp. 207–238.
[7] Jessica Dorsey, "The Erosion of Human(e) Judgement in Targeting? Quantification Logics, AI-enabled Decision Support Systems and Proportionality Assessments in IHL", *International Review of the Red Cross*, Vol. 107, No. 930, 2025.
[8] Natassha Selvaraj, "8 Machine Learning Models Explained in 20 Minutes", *Datacamp*, 16 September 2022, available at: https://www.datacamp.com/blog/machine-learning-models-explained.

This development creates a jurisprudential challenge that extends beyond familiar debates about compliance and accountability. At its core, IHL is not merely a technical regime regulating the conduct of hostilities – it is a normative system grounded in assumptions about human agency, moral responsibility, and the intrinsic value of human life. These commitments are reflected throughout the law's foundational principles, including distinction, proportionality, precaution, and humanity.[9] The prevailing legal discourse often assumes that existing IHL principles can simply be extended to AI-enabled systems through procedural safeguards, attribution rules, or human-supervision requirements.[10] While such measures are undoubtedly important, they may not fully address the deeper challenge posed by algorithmic mediation. Human decision-makers are increasingly required to rely upon outputs generated by systems whose internal reasoning processes may be difficult to understand, reconstruct, or meaningfully challenge. The concern is not that machines have become legal actors in their own right, nor that human decision-makers have disappeared from the battlefield[11] – rather, it is that algorithmic systems may progressively reshape how information is generated, interpreted, prioritised, and acted upon, thereby altering the conditions under which human judgment itself is exercised.

This article argues that preserving the legitimacy and vitality of IHL in an era of algorithmic warfare requires a reconceptualisation of the relationship between law, technology, and human dignity. Rather than treating dignity as a value protected only through *post hoc* accountability or judicial review, the article proposes that dignity should be understood as a design requirement shaping the architecture, governance, and deployment of military AI systems. This approach is conceptualised as "dignitary design". Dignitary design is not intended as a substitute for meaningful human control, nor does it replace existing obligations relating to weapons reviews, accountability, or compliance with the rules governing the conduct of hostilities. Instead, it provides a broader normative framework within which these mechanisms may be understood and developed. It proceeds from the premise that humanitarian values should not merely constrain the consequences of technological systems after deployment – they should be incorporated into the design, evaluation, governance, and operational use of those systems from the outset. In this sense, dignity becomes not only a legal value governing outcomes but also a structural principle shaping the conditions under which technological systems are developed and employed.

The urgency of this reframing arises from the growing gap between responsibility and agency in contemporary military operations. Traditional IHL

---

[9] For an overview of modern IHL, see Nils Melzer, *International Humanitarian Law: A Comprehensive Introduction*, ICRC, Geneva, 2019.

[10] Michael Pizzi, Mila Romanoff and Tim Engelhardt, "AI for Humanitarian Action: Human Rights and Ethics", *International Review of the Red Cross*, Vol. 102, No. 913, 2020.

[11] Aida Amanuel, "Removing the Human Decision-Making from Humanitarian Law An Examination of the Legal Implications of Autonomous Weapon Systems and Their Compliance with International Humanitarian Law", *Uppsala University*, Autumn Term 2024.

presumes that decisions to apply force are made by human actors capable of deliberation, contextual assessment, and moral evaluation, and legal accountability mechanisms are constructed accordingly: commanders assess proportionality, operators exercise precaution, legal advisers provide guidance, and States bear responsibility for internationally wrongful acts. AI complicates this architecture by distributing elements of cognition and assessment across complex socio-technical systems involving programmers, data curators, engineers, procurement authorities, military operators, commanders, and political institutions. As algorithmic systems increasingly contribute to operational judgments, responsibility becomes dispersed across multiple actors, while the practical locus of decision-making becomes more difficult to identify.[12]

This article does not advocate the prohibition of military AI, nor does it advance claims regarding machine consciousness or artificial legal personhood. Rather, it proceeds from a more modest but more demanding proposition: when States deploy systems capable of shaping decisions that may result in lethal effects, they bear an obligation to ensure that those systems are designed and governed in ways consistent with the normative commitments that underlie IHL. This requires moving beyond questions of compliance alone and considering how humanitarian values may be embedded within technological architectures themselves. From a jurisprudential perspective, such an approach represents an evolution rather than a departure from existing humanitarian law. Throughout its history, IHL has responded to technological change by reaffirming its underlying normative commitments while adapting their practical application to new means and methods of warfare. The principles of humanity and the dictates of public conscience, reflected most prominently in the Martens Clause,[13] provide a doctrinal foundation for extending humanitarian protections into domains not explicitly addressed by positive law. Artificial intelligence presents a particularly significant challenge because it affects not merely the means through which force is applied but the processes through which judgments concerning the use of force are formed. If IHL is to retain its human-centred identity under such conditions, dignity must be understood not only as a protected value but also as a principle capable of shaping technological design and governance.

The principal contribution of this article is therefore to reconceptualise human dignity as more than a moral aspiration or interpretive value within IHL. It argues that dignity should function as a design constraint that structures the development, deployment, and governance of military AI systems. In doing so, the article seeks to bridge existing debates concerning autonomy, accountability, and

---

[12] Sergey Sayapin, "International Law for a Fragile World: When Risk Outpaces Control – AI, Bioengineering, and Cyber-Autonomy", *Opinio Juris*, 23 April 2026, available at: https://opiniojuris.org/2026/04/23/international-law-for-a-fragile-world-when-risk-outpaces-control-ai-bioengineering-and-cyber-autonomy/.

[13] Antonio Cassese, "The Martens Clause: Half a Loaf or Simply Pie in the Sky", *European Journal of International Law*, Vol. 11, 2000.

human control while offering a framework capable of preserving the normative foundations of humanitarian law in an increasingly algorithmic battlespace.

The article develops this argument in seven sections. Section II examines the ontological status of artificial decision-making and its implications for legal agency. Section III situates human dignity within the normative structure of IHL, demonstrating its implicit centrality despite limited textual visibility. Section IV introduces the concept of dignitary design as a synthesis of law, ethics, and technology. Section V addresses accountability in hybrid human-machine systems and proposes expanded responsibility frameworks capable of addressing distributed agency. Section VI explores how dignity may be operationalised through proportionality assessments, precautionary obligations, weapons reviews, and algorithmic auditing. The conclusion reflects on the future legitimacy of IHL in a world where warfare is increasingly mediated by computational systems. Ultimately, the central question posed by military AI is not whether machines can comply with humanitarian law but whether humanitarian law can preserve its human-centred identity when operational judgment is increasingly shaped by machines. The answer depends upon whether dignity remains merely a rhetorical ideal or becomes a structural requirement of technological warfare.

II. The Ontological Problem: Can Machines "Decide" in Law?

The accelerating integration of artificial intelligence into military operations compels a foundational question that lies at the intersection of law, ethics, and technology: can machines meaningfully be said to "decide" in a juridical sense?[14] The answer bears directly on the integrity of IHL, which has traditionally assumed that legally significant acts in armed conflict originate in human deliberation and remain attributable to human actors. Without clarifying the ontological status of AI decision-making, contemporary debates concerning regulation, accountability, and compliance risk resting upon conceptual ambiguity.

At first glance, the language of "machine decision" appears metaphorical. Contemporary AI systems do not possess consciousness, intentionality, moral awareness, or the capacity for reflective judgment.[15] They process information through algorithmic architectures, generating outputs based on statistical inference, pattern recognition, and optimisation functions.[16] Even advanced machine-learning models do not currently deliberate in the manner associated with human practical

[14] Carlos Batallas, "When AI Meets the Laws of War", *Insights*, 3 October 2024, available at: https://www.ie.edu/insights/articles/when-ai-meets-the-laws-of-war/.
[15] Jeff Sebo and Robert Long, "Moral Considerations for AI Systems by 2030", *AI and Ethics*, Vol. 5, 2025.
[16] Antony Massobrio, "Machine Learning Optimization: Best Techniques and Algorithms", *Neural Concept*, 10 February 2023, available at: https://www.neuralconcept.com/post/ machine-learning-based-optimization-methods-use-cases-for-design-engineers.

reason.[17] Yet in operational settings, the outputs of such systems may shape target selection, collateral damage estimation, threat identification, force allocation, intelligence prioritisation, and other activities with profound legal consequences. When an algorithm classifies an object as a military objective, recommends engagement, or influences a proportionality assessment, the outcome is not metaphorical – it may contribute directly to decisions involving the use of lethal force.

At the same time, the challenge posed by AI should not be overstated. Military organisations have long relied upon technologies that mediate operational judgment.[18] Intelligence platforms, surveillance systems, predictive models, targeting software, and collateral damage estimation tools have influenced military decision-making for decades. The significance of contemporary AI therefore does not lie in the mere existence of technological mediation – rather, it lies in the increasing delegation of evaluative, predictive, and classificatory functions to systems capable of generating recommendations, prioritisation frameworks, and probabilistic assessments that shape human decision-making in ways that may be difficult to understand, contest, or reconstruct.[19] In other words, the novelty of AI resides less in automation as such than in the growing integration of machine-generated assessments into the informational architecture of military decision-making.

This distinction is important because not all forms of automation raise the same legal concerns. Conventional automated systems generally operate according to predefined rules and predictable parameters established by human designers.[20] By contrast, machine-learning systems may generate outputs through adaptive processes that are not fully transparent even to their developers.[21] The challenge posed by AI does not stem from the introduction of quantification into warfare. Quantitative reasoning has always formed part of military operations. Commanders routinely rely upon casualty estimates, intelligence assessments, risk matrices, operational models, and probabilistic forecasts. Proportionality analysis itself often incorporates predictive judgments regarding expected civilian harm and anticipated military advantage. The distinction between human and machine decision-making cannot

---

[17] J. M. Paterson, "AI Mimicking and Interpreting Humans: Legal and Ethical Reflections", *Journal of Bioethical Inquiry*, Vol. 22, No. 3, 2025.

[18] Patrick Hofstetter, Marius Geller and Florian Gerster, "The Impact of Artificial Intelligence On The Military Decision-Making Process and Mission Command", *The Defence Horizon Journal*, 11 September 2025, available at: https://tdhj.org/blog/post/ai-military-decision-making-2/.

[19] Tim Keating, Sarah Matar, Amr Fayad and Elena Chammas, "Preserving Human Moral Agency as use of AI-driven Autonomous Weapons Grows", *Kearney 100*, 15 May 2025, available at: https://www.kearney.com/industry/aerospace-defense/paris-air-show/artic le/preserving-human-moral-agency-as-use-of-ai-driven-autonomous-weapons-grows.

[20] Ajay Bandi, Bhavani Kongari, Roshini Naguru, Sahitya Pasnoor and Sri Vidya Vilipala, "The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges" *Future Internet*, Vol. 17, No. 9, 2025.

[21] Keonyoung Park and Ho Young Yoon, "AI Algorithm Transparency, Pipelines for trust not Prisms: Mitigating General Negative Attitudes and Enhancing Trust Toward AI", *Humanities and Social Sciences Communications*, Vol. 12, 2025.

therefore be reduced to a contrast between qualitative and quantitative reasoning. The more significant concern is whether algorithmic systems alter the relationship between calculation and judgment. When AI-generated outputs acquire an appearance of objectivity, precision, or authority, there is a risk that human decision-makers may defer to them without fully understanding their underlying assumptions, limitations, or uncertainties. The problem is not that machines perform calculations – rather, it is that machine-generated assessments may gradually become substitutes for independent human evaluation. Under such conditions, quantification risks evolving from a tool that informs judgment into a mechanism that structures or constrains it. The legal challenge therefore concerns not the existence of numerical reasoning but the possibility that computational outputs may increasingly shape the framework within which normative judgments are exercised.

One possible response is to insist that machines never truly decide and that responsibility therefore remains entirely with the human actors who deploy them.[22] Under this view, AI systems remain tools, however sophisticated, and all legally relevant decisions continue to originate with human operators, commanders, and States. This position possesses considerable intuitive appeal and remains broadly consistent with existing legal doctrine. Nevertheless, it risks becoming formalistic if applied without regard to the operational realities of contemporary military systems. In highly automated environments, human involvement may be temporally distant, informationally constrained, or structurally dependent upon algorithmic outputs. When commanders rely upon machine-generated assessments produced at speeds or levels of complexity beyond ordinary human cognitive capacities, their role may become increasingly supervisory rather than deliberative.[23]

A more accurate description of contemporary military AI is therefore neither replacement nor complete delegation, but hybridisation. Decision-making becomes distributed across complex socio-technical systems involving programmers, data curators, engineers, procurement authorities, legal advisers, commanders, operators, and algorithmic processes.[24] Programmers define architectures and optimisation objectives; data curators shape training datasets; engineers establish system parameters; legal advisers assess compliance; commanders determine operational goals; algorithms generate classifications and predictions; and operators execute or authorise actions. No single actor necessarily embodies the traditional image of autonomous judgment – instead, operational cognition becomes distributed across networks of human and technical actors. This phenomenon may be described as

[22] Robert Sparrow and Gene Flenady, "The Testimony Gap: Machines and Reasons", *Minds and Machines*, Vol. 35, 2025.
[23] Dmytro Shumlianskyi, "Some U.S. Military Personnel, Including Generals, Begin Using Generative Chatbots for Decision-Making", *Militarnyi*, 19 October 2025, available at: https://militarnyi.com/en/news/some-u-s-military-personnel-including-generals-begin-using-generative-chatbots-for-decision-making/.
[24] S. Sayapin, above note 12.

distributed agency.[25] The concept does not imply that machines possess legal personality or moral responsibility – rather, it recognises that legally significant outcomes increasingly emerge from interactions among multiple human and technical components. AI systems contribute to the production of operational knowledge but they do not become moral agents in their own right. Responsibility remains anchored in human and institutional actors even when the informational pathways leading to particular outcomes are shaped by algorithmic processes.[26]

This distinction is particularly important because some scholars have proposed forms of electronic personhood or artificial legal personality as responses to emerging technological challenges.[27] Such approaches do not resolve the concerns examined in this article. AI systems may improve situational awareness, reduce certain forms of human error, assist compliance with distinction obligations, and contribute to more informed operational decisions. The concern addressed here arises not from technological capacity itself but from the possibility that computational systems may reshape the conditions under which human judgment is exercised. The central question is therefore not whether machines can become moral agents – they cannot. Nor is it whether technology may assist compliance with IHL – it often can. The question is how legal systems can preserve meaningful human responsibility and moral judgment when operational cognition becomes increasingly distributed across human – machine assemblages.

This ontological challenge reveals the limitations of purely anthropocentric models of legal responsibility. Existing doctrines remain indispensable because responsibility must ultimately be borne by human and institutional actors. Yet those doctrines were developed in a context where legally significant decisions could generally be traced to identifiable individuals exercising relatively direct forms of judgment. AI complicates this picture by distributing cognition, knowledge production, and evaluative processes across complex socio-technical systems. The result is not the disappearance of human agency but its transformation. It is precisely this transformation that motivates the concept of dignitary design developed in the following sections. If artificial systems cannot internalise dignity as a moral value, dignity must instead be externalised as a normative constraint governing their design, deployment, and governance.[28] The task is therefore not to determine whether machines can become moral actors but to ensure that the technological systems through which military decisions are increasingly mediated remain structured by the humanitarian commitments that underlie IHL. The absence of machine morality thus becomes the reason for embedding human-centred safeguards into technological

---

[25] Werner Rammert, *Distributed Agency & Digital Technology: A Social Pragmatist View on Human – Technology Interaction*, Springer, Wiesbaden, 2025, pp. 129–142.

[26] Alexander Schwarz, "War Crimes", *Oxford Public International Law*, September 2014, available at: https://opil.ouplaw.com/display/10.1093/law:epil/9780199231690/law-9780199 231690-e431.

[27] See, for example, Pawe Nowik, "Electronic personhood for artificial intelligence in the workplace", *Computer Law & Security Review*, Vol. 42, 2021.

[28] See section IV below.

architectures themselves. Dignitary design thus emerges not as a technological aspiration but as a normative response to the realities of distributed agency in contemporary warfare.

III. Human Dignity as a Structural Principle of International Humanitarian Law

Human dignity occupies an unusual position within IHL. Unlike international human rights law (IHRL), where dignity is frequently articulated as an explicit foundational value,[29] IHL more rarely invokes the concept directly.[30] This relative textual silence, however, should not be mistaken for normative absence. While dignity may not appear as a comprehensive organising principle in the treaty language of IHL in the same manner as it does in many human rights instruments, it nevertheless performs a structural function within the humanitarian legal order. It informs the interpretation of core principles, underlies protections afforded to persons affected by armed conflict, and provides part of the normative foundation upon which the legitimacy of lawful warfare depends. Understanding dignity in this manner is essential for assessing how humanitarian protections may be preserved under conditions of increasing algorithmic mediation.

The historical development of IHL reflects a persistent effort to reconcile military necessity with respect for the humanity of persons affected by armed conflict.[31] From the earliest Geneva instruments of the nineteenth century to the contemporary law of armed conflict, humanitarian regulation has sought not merely to limit violence but to preserve a minimum moral order within situations characterised by organised force. This aspiration is evident in the gradual expansion of protections for the wounded and sick, prisoners of war, civilians, and other protected persons. While these rules are often expressed in operational terms, they embody a deeper normative commitment – the recognition that even during armed conflict human beings retain an irreducible moral status that constrains the permissible exercise of violence.

In this respect, dignity occupies a position analogous to that which it enjoys in IHRL, albeit through different doctrinal pathways. Human rights law frequently presents dignity as the foundation from which specific rights derive.[32] Humanitarian law proceeds differently – rather than establishing dignity as an explicit source of rights, it operationalises respect for human worth through principles governing the

---

[29] Sergey Sayapin, "The International Committee of the Red Cross and International Human Rights Law", *Human Rights Law Review*, Vol. 9, No. 1, 18 February 2009, pp. 95–126; Sergey Sayapin, "International Committee of the Red Cross and the Use of International Human Rights Law", in D. Rogers (ed.), *Human Rights in War*, Springer, Singapore, 2022, pp. 217–23 6.

[30] See Article 3(1)(c) common to the 1949 Geneva Conventions.

[31] Michael N. Schmitt, "Military Necessity and Humanity in International Humanitarian Law: Preserving the Delicate Balance", *Virginia Journal of International Law*, Vol. 50, No. 4, 2010.

[32] Oscar Schachter, "Human Dignity as a Normative Concept", *American Journal of International Law*, Vol. 77, No. 4, 1983.

conduct of hostilities and the treatment of protected persons. Yet the underlying normative concern remains comparable – both legal regimes reject the proposition that human beings may be reduced entirely to instruments of political, military, or strategic objectives. This anti-instrumentalisation function is particularly important. At its most fundamental level, dignity affirms that individuals possess value beyond their utility to collective goals.[33] IHL does not prohibit violence as such but it seeks to ensure that persons affected by warfare are not treated merely as objects through which military advantage is pursued. This idea is reflected throughout the structure of IHL: civilians may not be targeted because they are civilians; prisoners of war must be treated humanely regardless of their military significance; persons *hors de combat* retain legal protection despite their loss of operational utility. These obligations are difficult to explain solely through considerations of efficiency, reciprocity,[34] or military necessity. They derive their normative force from the recognition that persons retain moral significance even under conditions of armed conflict.

The clearest doctrinal expression of this orientation is found in the principle of humanity. Humanity has long occupied a central position within the humanitarian tradition and remains one of the foundational principles of IHL. Although its precise contours remain contested, the principle broadly requires that warfare be conducted within limits that respect the human condition of those affected by conflict.[35] It prohibits the infliction of suffering unnecessary for the achievement of legitimate military objectives and serves as a normative constraint upon the exercise of force. Humanity is therefore more than an ethical aspiration – it performs an interpretive and justificatory function within IHL, helping to explain why certain forms of conduct are prohibited even when they may offer military advantage.

Closely connected to the principle of humanity is the Martens Clause, first formulated in the 9th preambular paragraph of the 1899 Hague Convention (II) with Respect to the Laws and Customs of War on Land and subsequently reaffirmed in numerous humanitarian instruments, including Article 1(2) of Additional Protocol I. The Clause provides that, in situations not specifically addressed by treaty law, civilians and combatants remain under the protection of the principles of international law derived from established custom, the principles of humanity, and the dictates of public conscience. Obviously, the Martens Clause does not provide a complete legal regime governing military AI, nor does it resolve all questions concerning autonomy, accountability, or algorithmic decision-making. What it does provide is doctrinal support for the proposition that humanitarian interpretation must remain responsive to emerging forms of warfare even when positive law has

[33] Remy Debes, "Dignity", *The Stanford Encyclopedia of Philosophy*, 18 February 2023, available at: https://plato.stanford.edu/archives/spr2023/entries/dignity.
[34] ICRC, International Humanitarian Law Databases, Rule 140: Principle of Reciprocity, available at: https://ihl-databases.icrc.org/en/customary-ihl/v1/rule140.
[35] ICRC, "What is IHL?", 18 September 2015, available at: https://www.icrc.org/en/ document/what-ihl.

not yet developed detailed regulatory frameworks. In this respect, the Clause reinforces the broader argument advanced in this article – that technological innovation does not displace the normative commitments of IHL but instead requires their application under new conditions.

The relationship between dignity and positive humanitarian law becomes particularly visible in the treatment of protected persons. Common Article 3 of the Geneva Conventions prohibits, among other conduct, outrages upon personal dignity, particularly humiliating and degrading treatment. Similar concerns are reflected in provisions governing prisoners of war, civilians under occupation, and persons deprived of liberty. These rules do more than prohibit physical harm – they recognise that individuals possess interests in respect, recognition, and humane treatment that survive the onset of armed conflict. The influence of dignity extends beyond the treatment of protected persons and into the conduct of hostilities itself. Although the principles of distinction, proportionality, and precaution are generally expressed in operational language, they presuppose underlying judgments concerning the value of human life and the limits of permissible violence. [36] Distinction reflects the proposition that civilians may not be treated merely as obstacles to military success. Proportionality assumes that civilian lives possess significance that cannot be disregarded in pursuit of military objectives. Precaution requires parties to take feasible measures to minimise civilian harm, reflecting a continuing obligation to account for the welfare of those exposed to the effects of military operations. These principles are often analysed in technical terms, yet their normative coherence depends upon an underlying commitment to the worth of human beings.

This point deserves particular emphasis because proportionality is sometimes understood primarily as a balancing exercise between military advantage and civilian harm. Such a description is accurate but incomplete. Proportionality is not simply a mechanism for optimising military outcomes – its legitimacy depends upon the premise that civilian lives possess intrinsic value and cannot be treated merely as variables within strategic calculations. Without this premise, proportionality risks collapsing into a purely utilitarian exercise in which human beings are reduced to components of military cost-benefit analysis. Dignity supplies part of the normative foundation that prevents such a collapse. It reminds decision-makers that civilian harm is not merely an operational variable but a matter of moral and legal significance.

The same observation applies to precaution. The obligation to take feasible precautions reflects more than a desire to improve operational efficiency – it embodies a continuing concern for the protection of persons who may be affected by military operations. Precaution therefore illustrates how dignity operates not only as a retrospective standard for evaluating conduct but also as a prospective principle

[36] See section VI below.

shaping decision-making before harm occurs. In this respect, dignity influences not only outcomes but also processes.

Yet the structural role of dignity becomes more difficult to sustain when decision-making is increasingly mediated by algorithmic systems. Human dignity resists complete quantification. The value of a civilian life cannot be fully captured through statistical models, confidence scores, risk matrices, or optimisation functions. While such tools may provide useful information, they do not exhaust the normative considerations that underlie humanitarian decision-making. The challenge posed by AI therefore does not arise because dignity suddenly becomes relevant – it arises because dignity is already embedded within the normative architecture of IHL, while the technologies increasingly used to support military decision-making operate primarily through forms of classification, prediction, and optimisation. This tension reveals the deeper significance of dignity for contemporary debates concerning military AI.[37] If dignity functions as a structural principle of IHL, then the challenge is not simply ensuring that algorithmic systems comply with specific legal rules. It is ensuring that the normative commitments those rules embody continue to influence decision-making processes under conditions of technological mediation. In other words, the question is no longer whether dignity remains relevant to humanitarian law.[38] It is how dignity can continue to shape operational judgment when that judgment is increasingly informed, structured, and sometimes constrained by computational systems.

The answer cannot lie solely in *post hoc* accountability. Once dignity is recognised as a structural principle of IHL, it becomes necessary to consider how respect for human worth may be incorporated into the architecture of military technologies themselves. The challenge is therefore not merely interpretive but institutional and technological. This insight provides the foundation for the concept of dignitary design developed in the following section. If dignity already forms part of the normative architecture of humanitarian law, then preserving the human-centred character of IHL in an era of algorithmic warfare requires ensuring that dignity is reflected not only in legal evaluation after decisions are made but also in the systems through which those decisions are increasingly mediated.[39]

---

[37] Brian Stauffer, "A Hazard to Human Rights: Autonomous Weapons Systems and Digital Decision-Making", *Human Rights Watch*, 28 April 2025, available at: https://www.hrw.org/report/2025/04/28/a-hazard-to-human-rights/autonomous-weapons-systems-and-digital-decision-making.

[38] Ginevra Le Moli, "AI vs Human Dignity: When Human Underperformance is Legally Required", *Groupe d'études géopolitiques*, available at: https://geopolitique.eu/en/ articles/ai-vs-human-dignity-when-human-underperformance-is-legally-required/.

[39] ICRC, "Artificial Intelligence and Machine Learning in Armed Conflict: A Human-Centred Approach", *International Review of the Red Cross*, Vol. 102, No. 913, 2020, pp. 463–479.

IV. From Norm to Architecture: The Concept of Dignitary Design

The preceding sections have argued that military AI presents a challenge not because machines become moral agents but because algorithmic systems increasingly mediate the conditions under which human judgment is exercised. They have also argued that human dignity functions as a structural principle of IHL, informing the interpretation and application of core humanitarian norms even where it is not expressly articulated. The question that follows is therefore practical as well as jurisprudential: how can the normative commitments embodied in dignity continue to shape military decision-making when operational cognition is increasingly distributed across human – machine assemblages?

Existing legal and policy debates offer several partial answers: meaningful human control seeks to preserve human involvement in decisions concerning the use of force;[40] Article 36 weapons reviews provide mechanisms for assessing the legality of new weapons, means, and methods of warfare;[41] responsible AI frameworks emphasise transparency, explainability, fairness, and accountability; [42] and institutional accountability mechanisms seek to ensure that legal responsibility remains attributable to human actors and States. [43] Each of these approaches addresses an important dimension of the problem – yet none of them fully explains why preserving human-centred decision-making matters from the perspective of humanitarian law:

- meaningful human control, for example, identifies a procedural safeguard but does not itself provide a comprehensive normative account of what that safeguard is intended to protect;
- weapons reviews evaluate legality but do not necessarily address the broader humanitarian considerations embedded within technological systems;
- accountability frameworks determine who may be responsible when harm occurs but do not always guide how systems should be designed before deployment;
- responsible AI principles often emerge from broader governance debates and may not be tailored to the distinctive normative structure of IHL. As a result, existing approaches frequently operate

[40] Masoud Zamani, "How Meaningful is 'Meaningful Human Control' in LAWS Regulation?", *Articles of War*, 26 March 2025, available at: https://lieber.westpoint.edu/ how-meaningful-is-meaningful-human-control-laws-regulation/.

[41] *A Guide to the Legal Review of New Weapons, Means and Methods of Warfare: Measures to Implement Article 36 of Additional Protocol I of 1977*, ICRC, 2006, available at: https://library.icrc.org/library/docs/DOC/icrc_002_0902.pdf.

[42] Emmanouil Papagiannidis, Patrick Mikalef and Kieran Conboy, "Responsible Artificial Intelligence Governance: A Review and Research Framework", *The Journal of Strategic Information Systems*, Vol. 34, No. 2, 2025.

[43] Claudio Novelli, Mariarosaria Taddeo and Luciano Floridi, "Accountability in Artificial Intelligence: What it is and how it works", *AI & Society*, Vol. 39, 2023.

> as discrete mechanisms rather than as components of a unified humanitarian framework.

This article proposes dignitary design as such a framework. Dignitary design may be defined as the integration of respect for human dignity into the architecture, governance, evaluation, and operational deployment of military AI systems, with the objective of ensuring that technological mediation does not undermine the human-centred normative commitments of IHL. It seeks to preserve the role of dignity not merely as a value protected through legal review after harm has occurred but as a principle shaping the conditions under which military decisions are generated, informed, and implemented. Importantly, dignitary design does not require machines to recognise, understand, or internalise dignity – AI systems lack the moral capacities necessary for such recognition. They do not possess conscience, empathy, or the ability to appreciate the intrinsic worth of human beings. The objective of dignitary design is therefore not to create morally aware machines – rather, it is to ensure that human institutions embed dignity-sensitive constraints within the technological systems they develop and deploy. Precisely because machines cannot understand dignity, responsibility for preserving dignity must remain with the human actors and institutions that design, govern, and employ those systems.

Dignitary design thus differs from approaches that focus exclusively on outcomes. Traditional accountability frameworks often evaluate whether particular harms can be attributed to identifiable actors after an event has occurred, whereas dignitary design shifts attention upstream. It asks how technological systems should be structured before deployment, in order to preserve the humanitarian values that IHL seeks to protect. In this respect, it resembles the broader movement toward values-based design in technology governance,[44] while remaining grounded specifically in the normative commitments of humanitarian law. At its core, dignitary design rests upon a simple proposition: if dignity constitutes a structural principle of IHL, then military AI systems should be designed, governed, and deployed in ways that reflect that principle. The practical implications of this proposition may be organised into three interrelated dimensions – architectural, procedural, and institutional.

A. Architectural Dimension

The architectural dimension concerns the design characteristics of AI systems themselves. It focuses on how technological systems generate, present, and structure information relevant to military decision-making. If dignity requires that persons not be reduced to mere objects of optimisation, then the architecture of military AI systems should avoid treating civilian harm as a purely technical variable detached from its humanitarian significance. One implication is the importance of explainability. Human operators must be able to understand, at least to a meaningful

[44] See Gjalt de Graaf, *Value-Based Governance*, Cambridge University Press, Cambridge, 2025.

degree, the basis upon which significant recommendations or classifications are generated.[45] Explainability does not require complete transparency in every technical respect, nor does it imply that complex machine-learning models can always be rendered fully interpretable. It does, however, require that decision-makers possess sufficient information to evaluate, challenge, and contextualise algorithmic outputs. Systems that generate recommendations without providing meaningful opportunities for scrutiny risk transforming human oversight into a procedural formality.

The architectural dimension also encompasses the management of uncertainty.[46] Military AI systems frequently operate in environments characterised by incomplete information, rapidly changing conditions, and ambiguous data. Dignitary design therefore requires that uncertainty be made visible rather than concealed – that is, confidence scores, probability estimates, and margins of error should be presented in ways that encourage critical evaluation rather than passive acceptance. Human operators should be reminded not only of what a system predicts but also of the limitations of those predictions.

Similarly, system design should incorporate mechanisms for identifying and mitigating bias. [47] Training datasets, classification models, and optimisation objectives may systematically distort operational assessments in ways that increase the risk of civilian harm. Because dignity requires that individuals not be treated as expendable variables within computational processes, system architectures should be subject to ongoing testing, validation, and auditing designed to identify discriminatory or systematically harmful patterns.

Finally, architectural design should preserve meaningful opportunities for human intervention.[48] Override mechanisms, escalation procedures, and safeguards against fully automated implementation of high-consequence recommendations all serve to reinforce the continuing role of human judgment. The objective here is not to prevent technological assistance but to ensure that human actors retain genuine capacity to evaluate and, where necessary, reject algorithmic outputs.

---

[45] Markus Christen, Thomas Burri, Serhiy Kandul and Pascal Vörös, "Who is Controlling Whom? Reframing "Meaningful Human Control" of AI systems in Security", *Ethics and Information Technology*, Vol. 25, 2023.

[46] Jonathan Kwik, "Iterative Assessment for Military Artificial Intelligence (AI) Systems", in Bérénice Boutin, Taylor Kate Woodcock and Sadjad Soltanzadeh (eds), *Legal, Ethical, and Technical Dilemmas in Military Artificial Intelligence*, TMC Asser Press / Springer, 2026.

[47] Laura Bruun and Marta Bo, *Bias in Military Artificial Intelligence and Compliance with International Humanitarian Law*, Stockholm International Peace Research Institute, August 2025.

[48] "The Risks and Opportunities of AI on Humanitarian Action", *Wilton Park*, 15–17 May 2024, available at: https://www.wiltonpark.org.uk/reports/the-risks-and-opportunities-of-ai-on-humanitarian-action-report/.

B. Procedural Dimension

The procedural dimension concerns the processes through which military AI systems are evaluated, reviewed, tested, and monitored throughout their operational lifecycle. Existing legal obligations already provide important foundations for such processes. Most notably, Article 36 of Additional Protocol I requires States to determine whether the employment of new weapons, means, or methods of warfare would be prohibited under international law. Yet the emergence of adaptive and data-driven systems raises questions that extend beyond traditional weapons reviews. Unlike many conventional weapons, AI-enabled systems may evolve over time through software updates, retraining processes, changing datasets, and new operational environments.[49] Their behaviour may therefore depend not only upon their original design but also upon subsequent interactions with data and users. Dignitary design responds to this reality by emphasising lifecycle governance rather than one-time review.

Under a dignitary design framework, legal review becomes an ongoing process rather than a singular event. Systems should be subject to continuous assessment before, during, and after deployment. Testing should examine not only technical performance but also the ways in which algorithmic outputs influence human decision-making. Auditing mechanisms should evaluate whether systems systematically encourage excessive reliance upon automated recommendations or diminish opportunities for meaningful human evaluation.[50] In this respect, dignitary design extends rather than replaces existing review mechanisms. Article 36 reviews, legal assessments, operational testing, and responsible AI evaluations remain indispensable. The contribution of dignitary design lies in providing a normative lens through which these processes may be understood and integrated.

C. Institutional Dimension

The institutional dimension concerns the allocation of responsibilities among the actors who shape military AI systems. The concept of distributed agency developed in section II suggests that legally significant outcomes increasingly emerge from interactions among multiple human and technical components. Preserving dignity under such conditions requires institutional arrangements capable of addressing this complexity.

Responsibility cannot rest exclusively with commanders or operators: developers influence system capabilities through design choices; data curators shape the informational foundations upon which systems learn and operate; procurement authorities determine which technologies enter operational environments; legal

---

[49] "AI Model Retraining", *Transforming Data with Intelligence (TDWI)*, available at: https://tdwi.org/pages/glossary/ai-model-retraining.aspx.

[50] "Independent AI Audit Requirements", *VerifyWise*, available at: https://verifywise.ai/lexicon/independent-ai-audit-requirements.

advisers assess compliance and identify risks; and political authorities establish strategic priorities and governance frameworks. Each of these actors contributes to the conditions under which algorithmic systems generate operational consequences. Legal advisers occupy a notable position within this framework.[51] Institutional responsibility also requires transparency and oversight mechanisms capable of generating organisational learning. Post-strike assessments,[52] incident investigations, and review procedures should not focus exclusively on assigning blame after failures occur. They should also identify structural weaknesses within socio-technical systems and facilitate improvements designed to reduce future risks.

D. Dignity as a Design Constraint

Taken together, these architectural, procedural, and institutional dimensions reveal the distinctive contribution of dignitary design. Meaningful human control, Article 36 review, explainability, auditing, and accountability remain indispensable mechanisms. Dignitary design does not replace them – rather, it provides the normative framework that explains why they matter and how they should be oriented. The central insight is that dignity should function as a design constraint rather than merely a retrospective evaluative standard. Traditionally, dignity enters legal analysis after harm has occurred, informing assessments of responsibility, legality, or reparation. Dignitary design shifts attention to an earlier stage – it asks whether technological systems are structured in ways that preserve the status of human beings as subjects of legal concern rather than objects of computational optimisation. This shift has significant implications for the future development of military AI. It suggests that compliance with IHL cannot be evaluated solely by examining outcomes – the processes through which decisions are generated also matter. A system that consistently obscures uncertainty, discourages scrutiny, or transforms civilian lives into abstract variables may undermine humanitarian values even if particular outcomes remain formally lawful. Conversely, systems designed to facilitate critical evaluation, preserve meaningful human judgment, and foreground humanitarian considerations may strengthen the capacity of military institutions to comply with IHL under increasingly complex conditions.

Dignitary design therefore represents neither a technological solution nor a new legal doctrine. It is a jurisprudential framework that seeks to align technological development with the normative commitments that have long animated IHL. By embedding dignity within the architecture, governance, and deployment of military AI systems, it offers a means of preserving the human-centred identity of IHL in an era of algorithmic warfare.

---

[51] See section V below.

[52] Jan Frelin, *Purposeful Assessment: Assessing the Effects of Military Operations*, FOI, Swedish Defence Research Agency, Stockholm, January 2009.

V. Accountability in Hybrid Human – Machine Assemblages

The emergence of military AI is frequently presented as an accountability problem. As algorithmic systems assume increasingly significant roles in targeting, surveillance, intelligence analysis, and operational planning, concerns arise regarding who should bear responsibility when harmful outcomes occur.[53] These concerns are understandable – IHL has traditionally been premised upon the attribution of legally significant acts to identifiable human actors and States. Yet contemporary military AI complicates this picture by distributing elements of cognition, assessment, and decision-support across complex socio-technical systems. The challenge is therefore not merely one of legal attribution – it is a question of how responsibility should be structured in operational environments characterised by distributed agency. The accountability challenge therefore lies not in identifying a single replacement for the traditional decision-maker but in understanding how responsibility should be allocated within increasingly complex networks of human and technical actors.

This perspective requires moving beyond simplistic narratives suggesting that AI creates accountability vacuums. The notion of an accountability gap has become a recurring theme within debates concerning autonomous systems.[54] According to this view, harmful outcomes may become difficult to attribute because no individual actor exercises complete control over the processes leading to a particular decision. While such concerns should not be dismissed, the language of accountability gaps may be misleading. Military operations have rarely been the product of isolated individual choices – even before the emergence of contemporary AI, decisions concerning the use of force depended upon multiple actors, institutions, information sources, and chains of authority. What AI does is not create complexity from nothing – rather, it renders existing complexity more visible while introducing additional layers of technological mediation. This observation suggests that accountability in military AI should be understood not primarily as a problem of absence but as a problem of distribution. In this respect, military AI may be understood through the lens of the "many-hands problem" – a concept developed in fields such as technology ethics, public administration, and risk governance.[55] The many-hands problem arises when harm results from the actions of multiple actors, none of whom individually controls the entire process leading to a particular outcome. Under such conditions, responsibility may become fragmented, diluted, or obscured even though numerous actors contribute to the relevant chain of events.

---

[53] Nuno Marques, "Artificial Intelligence, Causation and State Responsibility in Cyberspace", *International Law Blog*, 23 March 2026, available at: https://internationallaw.blog/2026/03/23/artificial-intelligence-causation-and-state-responsibility-in-cyberspace/.

[54] Peter Königs, "Artificial Intelligence and Responsibility Gaps: What is the Problem?", *Ethics and Information Technology*, Vol. 24, 2022.

[55] Matthijs M. Maas, *Architectures of Global AI Governance: From Technological Change to Human Choice*, Oxford University Press, Oxford, 2025, pp. 497–593.

Military AI presents a particularly clear example of this phenomenon: developers influence system behaviour through design choices and optimisation objectives;[56] data curators shape training datasets that affect classification accuracy and predictive performance; procurement authorities determine which technologies are acquired and deployed; legal advisers assess compliance with humanitarian law; commanders establish operational objectives and rules of engagement;[57] operators decide whether and how to act upon algorithmic recommendations; and political authorities define broader strategic priorities. Each actor contributes to the operational environment within which AI systems generate legally significant effects – yet no single actor necessarily exercises complete control over the entire process. Recognising the distributed nature of responsibility does not require abandoning traditional accountability frameworks. On the contrary, existing doctrines remain indispensable – individual criminal responsibility,[58] command responsibility,[59] and State responsibility continue to provide essential mechanisms for enforcing compliance with IHL.[60] However, these doctrines were largely developed in contexts where legally significant conduct could be more readily linked to identifiable human decisions.[61] AI does not render these frameworks obsolete but it does expose some of their limitations when applied to highly distributed decision-making environments.

The limitations of traditional accountability models become particularly apparent when responsibility is conceptualised exclusively in terms of individual fault. Consider a military AI system that incorrectly classifies a civilian structure as a military objective.[62] The resulting harm may reflect a combination of factors – flaws in training data, design assumptions, operational pressures, inadequate oversight, inappropriate deployment conditions, or excessive reliance upon algorithmic recommendations. Focusing exclusively upon the final operator risks obscuring the broader institutional and technological context within which the error occurred. Conversely, dispersing responsibility too broadly risks creating a situation in which no actor is meaningfully accountable. The challenge is therefore to develop

---

[56] Hayden Horner, "Engineers Navigating Ethics of AI Development in Line with UNESCO Principles", *Engineering Institute of Technology*, 3 June 2025, available at: https://www.eit.edu.au/engineers-ethics-ai-development-unesco-principles/.

[57] Ann-Katrien Oimann and Adriana Salatino, "Command Responsibility in Military AI Contexts: Balancing Theory and Practicality", *AI Ethics*, Vol. 5, 2025.

[58] See Ciara Damgaard, *Individual Criminal Responsibility for Core International Crimes: Selected Pertinent Issues*, Springer, Berlin / Heidelberg, 2008.

[59] Yoram Dinstein, "Command Responsibility", *Oxford Public International Law*, September 2015, available at: https://opil.ouplaw.com/display/10.1093/law:epil/97801992 31690/law-9780199231690-e273.

[60] Sergey Sayapin, *A Central Asian Perspective on International Law*, Hart Publishing, 2026, p. 219.

[61] Sergey Sayapin, "International Law for a Fragile World: Global Risk and the Limits of Classical International Law", *Opinio Juris,* available at: https://opiniojuris.org/2026/ 02/06/international-law-for-a-fragile-world-global-risk-and-the-limits-of-classical-international-law/.

[62] Jimena Sofia Viveros Álvarez, "The Risks and Inefficacies of AI systems in Military Targeting Support", *Humanitarian Law & Policy*, 4 September 2024, available at: https:// blogs.icrc.org/law-and-policy/2024/09/04/the-risks-and-inefficacies-of-ai-systems-in-military-targeting-support/.

accountability structures capable of recognising distributed causation without allowing responsibility to dissipate. This challenge is particularly significant because AI systems lack the moral capacities necessary for legal responsibility. Machines do not possess intentions, understand legal obligations, or appreciate the consequences of their actions in the manner required for criminal or moral accountability – consequently, proposals to attribute responsibility directly to AI systems provide little assistance in addressing humanitarian concerns. The problem confronting IHL is not the absence of an appropriate legal category for machines – it is ensuring that responsibility remains meaningfully connected to the human and institutional actors who design, govern, and deploy technological systems.

For this reason, State responsibility assumes heightened importance in the context of military AI. Under international law, States remain responsible for internationally wrongful acts attributable to them, regardless of whether those acts involve advanced technologies or traditional means of warfare.[63] The deployment of AI systems does not diminish these obligations. If anything, it reinforces the need for robust institutional safeguards capable of ensuring compliance with humanitarian norms. States cannot avoid responsibility by pointing to algorithmic complexity, software errors, or technological opacity – the use of AI therefore strengthens rather than weakens the importance of institutional due diligence.[64] The latter may be understood as the obligation to establish governance structures capable of identifying, managing, and mitigating risks associated with military AI. Such obligations extend beyond the narrow question of whether a particular system is technically functional – they encompass procurement procedures, testing protocols, legal review mechanisms, auditing processes, training requirements, operational oversight, and post-deployment evaluation. From the perspective of dignitary design, due diligence serves as a means of ensuring that humanitarian considerations remain embedded throughout the lifecycle of military AI systems.

This broader conception of responsibility also highlights the importance of legal advisers.[65] Article 82 of Additional Protocol I requires States to ensure that legal advisers are available to advise military commanders on the application of IHL. Traditionally, this function has focused on operational compliance and legal review. In the context of military AI, however, legal advisers assume a more expansive role. Their expertise may contribute not only to the assessment of particular military operations but also to the development of governance frameworks governing technological systems themselves. They may participate in weapons reviews, contribute to procurement decisions, evaluate testing procedures, advise on operational safeguards, assist in developing auditing mechanisms, and help identify risks associated with algorithmic decision-making. Their role is therefore not limited

[63] See ARSIWA, Article 1.
[64] Beyza Unal and Ulysse Richard, *Governance of Artificial Intelligence in the Military Domain*, UNODA Occasional Papers, No. 42, United Nations, New York, June 2024, pp. vi, 23.
[65] Leslie C. Green, "The Role of Legal Advisers in the Armed Forces", *Israel Yearbook on Human Rights*, Vol. 7, 1977.

to determining whether a system satisfies minimum legal requirements – it includes helping institutions understand how technological choices may affect the human-centred values embedded within IHL.

Notably, the importance of institutional governance extends beyond legal advisers alone – effective accountability requires organisational structures capable of generating transparency, oversight, and learning.[66] Dignitary design provides a useful framework for understanding these responsibilities because it links accountability to the preservation of human dignity rather than merely to the allocation of blame. Institutions are therefore responsible not merely for preventing unlawful outcomes but for ensuring that technological systems are governed in ways consistent with the humanitarian values that underlie IHL.

Ultimately, the challenge posed by military AI is not that responsibility disappears – it is that responsibility becomes more difficult to organise and more important to maintain. The preservation of dignity under conditions of distributed agency therefore depends upon governance structures capable of ensuring that responsibility remains visible throughout the lifecycle of military AI systems. If section IV demonstrated that dignity should function as a design constraint, the next question becomes how that constraint may be operationalised within specific decision-making processes. It is to that question – particularly in relation to proportionality, precaution, auditing, and operational review – that the article now turns.

VI. Operationalising Dignity: Proportionality, Precaution, and Algorithmic War

If dignity is to function as a meaningful constraint on military AI rather than a purely aspirational value, it must be translated into operational practices capable of guiding decision-making in concrete contexts. The challenge is not simply one of legal interpretation – it is a question of institutional design, operational governance, and technological architecture. The preceding sections have argued that dignity functions as a structural principle of IHL and that military AI increasingly mediates the conditions under which humanitarian judgments are exercised. The practical question that follows is therefore how dignity may continue to influence decision-making processes when those processes increasingly incorporate algorithmic systems.

A. Dignity and Proportionality in Algorithmic Warfare

The principle of proportionality provides a useful illustration of the challenges posed by AI-mediated decision-making. Proportionality prohibits attacks expected to cause incidental civilian harm that would be excessive in relation to the concrete and direct

[66] Daniel E. Esser and Heiner Janus, "How are Accountability and Organisational Learning Related? A Study of German Bilateral Development Cooperation", *Evaluation*, Vol. 29, No. 4, 2023.

military advantage anticipated.[67] The application of this principle necessarily requires predictive judgments under conditions of uncertainty. Commanders routinely assess intelligence, estimate potential civilian casualties, evaluate operational alternatives, and weigh anticipated military benefits – quantification has therefore always formed part of proportionality analysis. Contemporary military decision-making already relies upon casualty estimates, collateral damage estimation tools, risk assessments, and probabilistic judgments.[68] The challenge posed by AI lies elsewhere – it arises when algorithmic outputs begin to shape proportionality assessments in ways that obscure uncertainty, discourage independent evaluation, or transform predictive calculations into *de facto* decision rules. AI systems may contribute valuable information to proportionality assessments by identifying patterns, modelling potential effects, or processing large volumes of operational data.[69] Yet proportionality remains fundamentally a normative judgment – no algorithm can ever determine whether anticipated civilian harm is legally "excessive" in relation to expected military advantage without relying upon assumptions, proxies, and evaluative choices embedded within its design. The risk is therefore not that AI introduces quantification into proportionality analysis but that quantitative outputs may acquire an appearance of objectivity that discourages critical scrutiny.

From the perspective of dignitary design, proportionality assessments should preserve meaningful opportunities for human judgment. Human decision-makers should be able to understand how relevant outputs were generated, identify sources of uncertainty, evaluate alternative courses of action, and challenge algorithmic recommendations where appropriate. Dignity is thus protected not by excluding quantitative tools from military planning but by ensuring that quantitative assessments remain embedded within a broader framework of humanitarian evaluation. This concern is particularly important because dignity functions, in part, as a safeguard against the reduction of human beings to variables within optimisation processes – civilian lives cannot be understood solely as data points within a cost-benefit calculation. While proportionality necessarily involves balancing, the legitimacy of that balancing exercise depends on recognising the intrinsic value of those who may be harmed – dignitary design requires that military AI systems avoid presenting civilian harm as a purely technical parameter detached from its humanitarian significance.

---

[67] ICRC, "The Principle of Proportionality", available at: https://www.icrc.org/sites/default/files/wysiwyg/war-and-law/04_proportionality-0.pdf.

[68] Zena Assaad, "Data, Decisions, and Death: Examining the Role of AI in Contemporary Warfare", *Arab Reform Initiative*, 27 May 2026, available at: https://www.arab-reform .net/publication/data-decisions-and-death-examining-ai-in-contemporary-warfare/.

[69] "Key Takeaways of The Military AI, Peace & Security Dialogues 2025", *UNODA*, 17 September 2025, available at: https://disarmament.unoda.org/en/updates/key-takeaways-military-ai-peace-security-dialogues-2025.

B. Precaution and the Governance of Technological Risk

The principle of precaution provides a second important avenue through which dignity may be operationalised.[70] Unlike proportionality, which focuses on balancing competing considerations in particular operational contexts, precaution is inherently future-oriented. It requires parties to take feasible measures designed to minimise civilian harm before attacks occur. In this respect, precaution is especially relevant to the governance of emerging technologies.

Military AI introduces new forms of uncertainty that may not be fully captured by traditional operational planning – as discussed above, AI performance may vary across environments, datasets may contain hidden biases, system behaviour may change following updates, and operators may develop inappropriate reliance upon algorithmic recommendations. These risks cannot always be addressed through retrospective accountability mechanisms, and require anticipatory governance – dignitary design therefore encourages a precautionary approach to military AI. In particular, it requires institutions to identify foreseeable risks and implement measures designed to reduce them – testing, verification, validation, simulation, red-teaming,[71] scenario-based evaluation, and independent review all serve this function. These practices help ensure that systems are assessed not only for technical reliability but also for their potential effects on humanitarian decision-making. Significantly, the precautionary dimension of dignitary design also extends to operational deployment – restrictions on the use of particular systems, limitations on operational environments, requirements for human review, and mechanisms for suspending deployment where performance becomes unreliable may all constitute precautionary measures. Such safeguards recognise that human dignity may be threatened not only by unlawful outcomes but also by institutional failures to anticipate foreseeable risks.

C. Dignity-Based Auditing

One of the most practical implications of dignitary design is the development of dignity-based auditing mechanisms. Traditional evaluations of military technologies often focus on effectiveness, reliability, accuracy, and legality.[72] While these considerations remain essential, they may not fully capture the broader humanitarian implications of AI-mediated decision-making. A dignity-based audit should ask not only whether a system functions as intended but also, importantly, whether it

---

[70] ICRC, International Humanitarian Law Databases, Rule 15: Principle of Precautions in Attack, available at: https://ihl-databases.icrc.org/en/customary-ihl/v1/rule15.

[71] "What is Red Teaming? Methodology and Scope of a Red Team Operation", *Vaadata*, 7 April 2026, available at: https://www.vaadata.com/en/blog/what-is-red-teaming-method ology-and-scope-of-a-red-team-operation/.

[72] Netta Goussac and Rain Liivoja, "Legal Reviews of Military Artificial Intelligence Capabilities", *Asia-Pacific Institute for Law and Security (APILS)*, 25 August 2025, available at: https://apils.org/2025/08/25/legal-reviews-of-military-artificial-intelligence-capabilities/.

preserves the conditions necessary for meaningful human judgment. Such an audit might examine, in particular, whether uncertainty is communicated transparently, whether algorithmic outputs can be challenged and reviewed, whether civilian protection is adequately represented within system design, and whether operators retain genuine opportunities to exercise independent evaluation. Similarly, dignity-based auditing should assess the extent to which systems incorporate safeguards against bias, overreliance, and inappropriate automation. It should also examine whether confidence scores are presented responsibly, whether alternative interpretations of data remain visible, and whether system architectures encourage critical engagement rather than passive acceptance of algorithmic recommendations.

D. Rules of Engagement and Operational Governance

Dignitary design also has implications for operational governance. Rules of engagement (ROE) have long served as mechanisms through which legal obligations are translated into operational guidance.[73] As AI systems become more deeply integrated into military decision-making, existing ROE frameworks may require adaptation to address new forms of technological mediation. AI-specific ROE might establish conditions governing when algorithmic recommendations may be relied upon, when additional human review is required, and when deployment should be suspended. Their rules could specify confidence thresholds below which recommendations cannot be acted upon without further verification. They might require escalation procedures where uncertainty exceeds predefined limits or where civilian harm cannot be assessed with sufficient confidence. Operational governance may also require restrictions on the contexts in which particular systems may be used – certain environments may present levels of complexity, uncertainty, or civilian presence that render heavy reliance on AI especially problematic. Dignitary design therefore supports context-sensitive governance rather than uniform assumptions regarding the suitability of particular technologies. More broadly, operational governance should preserve opportunities for human intervention throughout the decision-making process.

E. From Weapons Review to Lifecycle Governance

The operationalisation of dignity also requires reconsidering how military AI systems are reviewed over time. Existing legal obligations, particularly those associated with Article 36 of Additional Protocol I, provide essential mechanisms for evaluating new weapons, means, and methods of warfare – yet many AI systems differ from traditional military technologies in ways that complicate one-time review processes. As discussed above, software updates, retraining procedures, evolving datasets, and changing operational environments may alter system performance after deployment – a system assessed as reliable under one set of conditions may behave

---

[73] *The Sanremo Handbook on Rules of Engagement*, International Institute of Humanitarian Law, 2009, available at: https://iihl.org/wp-content/uploads/2022/12/ROE-HANDBOOK-ENGLISH.pdf.

differently under another. Dignitary design therefore supports the concept of lifecycle governance – rather than treating review as a singular event occurring prior to deployment, lifecycle governance recognises the need for continuous evaluation throughout the operational existence of a system. Monitoring, auditing, testing, incident review, and periodic reassessment should become essential components of humanitarian governance.[74] Such an approach does not replace Article 36 review – on the contrary, it builds upon its underlying logic: if States have an obligation to assess the legality of new military technologies before deployment, it is difficult to justify ignoring significant changes that may occur after deployment. Lifecycle governance therefore represents a natural extension of existing humanitarian obligations into the context of adaptive and data-driven systems.

VII. Conclusion

The emergence of artificial intelligence in military operations has generated intense debate concerning autonomy, accountability, and the future of warfare. Much of this discussion has focused on whether existing legal frameworks are capable of regulating increasingly sophisticated technologies. While these concerns remain important, this article has argued that they do not fully capture the challenge posed by military AI. The central question is not whether machines can become legal or moral actors, nor whether technological systems can simply be accommodated within existing regulatory structures. Rather, it is whether the human-centred normative commitments that underlie IHL can be preserved when the processes through which military decisions are generated become increasingly mediated by algorithmic systems.

The analysis has proceeded from the observation that contemporary military AI does not eliminate human agency but transforms the conditions under which agency is exercised. Decisions concerning the use of force increasingly emerge from interactions among commanders, operators, legal advisers, developers, procurement authorities, data curators, and algorithmic systems. Accordingly, operational cognition becomes distributed across complex socio-technical networks in which human actors remain formally responsible while relying upon informational structures that may be difficult to understand, challenge, or reconstruct. The challenge confronting IHL is therefore not the disappearance of responsibility but its reconfiguration under conditions of distributed agency.

At the same time, the article has sought to avoid overstating the novelty of technological change. Military decision-making has long relied upon technologies that mediate perception, communication, targeting, and operational planning. Artificial intelligence does not introduce technological mediation into warfare for the

[74] Alexandra T. Warner, *What is Monitoring in Humanitarian Action? Describing Practice and Identifying Challenges*, Active Learning Network for Accountability and Performance (ALNAP), London, 2017, pp. 11–14.

first time. Its significance lies in the increasing delegation of evaluative, predictive, and classificatory functions to systems that shape the informational environment within which humanitarian judgments are made. The legal challenge arises not because machines become moral actors but because algorithmic systems increasingly influence how human actors understand and evaluate the battlefield.

To address this challenge, the article has argued that human dignity should be understood as a structural principle of IHL. Although dignity occupies a less explicit position within humanitarian law than within international human rights law, it nevertheless informs the normative architecture of the humanitarian legal order. Through the principles of humanity, distinction, proportionality, precaution, and humane treatment, IHL affirms that persons affected by armed conflict retain an irreducible moral status that constrains the exercise of violence. Dignity functions, in part, as a safeguard against the instrumentalisation of human beings, ensuring that individuals are not reduced to mere objects through which military objectives are pursued. This understanding of dignity provides the foundation for the article's central contribution – the concept of dignitary design. Dignitary design is not proposed as a new regulatory mechanism, a substitute for meaningful human control, or an alternative to existing accountability frameworks. Rather, it is a jurisprudential framework for integrating respect for human dignity into the architecture, governance, evaluation, and operational deployment of military AI systems. It seeks to ensure that technological mediation does not erode the humanitarian commitments underlying IHL by embedding those commitments within the systems and institutions through which military decisions are increasingly shaped.

The article has further argued that preserving dignity in an era of military AI requires attention to three interconnected dimensions of governance. Architecturally, systems should be designed to support meaningful human judgment through explainability, transparency regarding uncertainty, safeguards against bias, and opportunities for human intervention. Procedurally, dignity requires ongoing review, auditing, testing, and lifecycle governance capable of addressing the evolving character of data-driven systems. Institutionally, it requires responsibility structures that reflect the distributed nature of contemporary military decision-making while ensuring that accountability remains anchored in human and institutional actors.

Operationally, dignitary design finds expression through proportionality assessments, precautionary measures, dignity-based auditing, rules of engagement, and continuous review mechanisms. These practices do not prohibit algorithmic assistance – rather, they seek to ensure that algorithmic systems remain subordinate to humanitarian judgment. The objective is not to prevent the use of AI in armed conflict but to preserve the capacity of military institutions to make decisions consistent with the values that humanitarian law exists to protect.

Ultimately, the challenge posed by military AI extends beyond questions of technological regulation – it concerns the future normative identity of IHL itself. Throughout its history, humanitarian law has adapted to changing means and methods of warfare while preserving its underlying commitment to humanity. Artificial intelligence presents a particularly significant test because it affects not only the instruments through which force is applied but also the processes through which judgments concerning force are formed. If humanitarian law is to remain a genuinely human-centred legal regime under such conditions, dignity cannot be treated solely as a retrospective standard for evaluating outcomes. It must also function as a principle shaping the systems through which decisions are generated.

The future of IHL will therefore depend not only on whether new technologies can be accommodated within existing legal frameworks but also on whether those frameworks continue to preserve human beings whose protection constitutes their fundamental purpose. In an increasingly algorithmic battlespace, the preservation of human dignity is not a peripheral ethical concern. It is a condition of maintaining legitimacy, coherence, and humanitarian character of the law itself. Dignitary design offers one possible path toward that objective by ensuring that, even as warfare becomes more technologically sophisticated, the normative centre of humanitarian law remains firmly human.